Dealing with Execution Uncertainty in the Continuous Double Auction

Gert van Valkenhoef

September 2008 - April 2009

Master Thesis
Artificial Intelligence
Department of Artificial Intelligence
University of Groningen, The Netherlands

Internal Supervisor:
Prof. dr. Rineke Verbrugge (University of Groningen, The Netherlands)

External Supervisor:
Prof. dr. Nick R. Jennings (University of Southampton, United Kingdom)

Abstract

With the advent of Grid computing, peer-to-peer systems and ad-hoc networks, distributed systems are becoming increasingly open and operate at a much larger scale in terms of the number of computational entities in the system. This poses new challenges for computational resource allocation. Specifically, a resource allocation mechanism must now deal with the increased size of the allocation problem, as well as the fact that computational entities may be controlled by different parties, with conflicting interests. Finally, there is an increased risk that a computational entity will fail to perform its assigned tasks: there is *execution uncertainty*. Traditional resource allocation mechanisms are inadequate in this setting. Thus, a new mechanism is required.

To this end, in this work an agent-based system is developed that can solve the resource allocation problem in large-scale, open, distributed systems. Specifically, we develop the Trust-Based CDA (T-CDA), a decentralised market-based resource allocation mechanism. In this system, computational entities are modelled as agents that may buy or sell the use of resources. The T-CDA is an extension of the Continuous Double Auction (CDA) mechanism, which is a decentralised market-based resource allocation system. This means that the CDA can deal with the first two challenges of large-scale, open, distributed systems: its decentralised nature allows it to deal with large allocation problems, while market-based mechanisms can deal with different agents having conflicting interests. Now, to meet the execution uncertainty challenge, the T-CDA additionally allows agents to use a *trust model* in deciding whether or not to trade with a certain other agent. Specifically, an additional step is introduced in the trading process that allows agents commit to trades they believe will maximise their expected utility.

We empirically evaluate the mechanism with Zero-Intelligence (ZI) agents, both against the optimal solution given complete and perfect information and against the standard CDA. We show the T-CDA consistently outperforms the traditional CDA as execution uncertainty increases in the system. Furthermore, we investigate the robustness of the mechanism to unreliable trust information and find that performance degrades gracefully as information quality decreases.

Now, because in a decentralised mechanism the individual agent behaviours are an important determinant in overall system behaviour, we also develop Trust-Based ZIP (T-ZIP), a rudimentary trading strategy for the T-CDA. The T-ZIP strategy is empirically evaluated and shown to outperform ZI in specific conditions, showing an increase of efficiency of up to 80%. However, the T-ZIP is shown to fail in other conditions and thus is not a general trading strategy for the T-CDA. Insights are provided into the failure of T-ZIP in these conditions and ways to design a generally applicable strategy are identified.

Acknowledgements

First of all, I would like to thank Professor Nick Jennings, my supervisor in Southampton, for his excellent advise and support throughout my thesis work. Most of all, I value his many intelligent, thorough and timely reviews of my written work, all done in spite of the vast army of PhD students vying for his attention. He has picked apart even the most carefully constructed pieces of text, vastly improving the way I write and read scientific texts.

Then, no less important, I wish to thank my supervisor at the University of Groningen, Dr. Rineke Verbrugge, for her enthusiastic support throughout my master's degree. I have had the pleasure of working with her not only during my thesis project: she has provided helpful advise throughout my master, she has supervised a minor project and I have been an assistant on a number of her courses. Without her, I would not have enrolled in the Artificial Intelligence master's programme altogether.

My thanks also go out to Dr. Sarvapali D. Ramchurn (Gopal) and Dr. Perukrishnen Vytelingum (Krishnen), both research fellows at the University of Southampton, for their day-to-day support of my work and the many interesting discussions that have no doubt helped me move in the right direction. In addition, I would like to thank them and the rest of the Agents group at Southampton for welcoming me in their midst and providing a pleasant working environment as well as the necessary distractions: lunch, coffee, the obligatory banter and the occasional night out.

I also thank my family and, specifically, my parents for their unwavering moral, practical and financial support throughout my studies. Last, but not least, I thank my amazing girlfriend, Brechtsje, for being there for me all the time, even though we were far apart.

Table of Contents

Chapter 1

Introduction

Resource allocation is an important problem in computer science. Traditionally, it has been studied in settings where computational entities are cooperative (i.e. they work together towards a shared goal, such as the completion of some computation) and the allocation is determined by a central authority (e.g. the operating system kernel allocating available CPU time to different processes, or a router allocating available network bandwidth to different services).

However, with the advent of Grid computing, peer-to-peer systems and adhoc networks, distributed systems are now being populated by an increasingly large number of computational entities. The following examples illustrate this trend:

- The current (April 2009) top-ranked supercomputer,[1] IBM Roadrunner, consists of 129.600 processor cores;
- The Large Hadron Collider Computing Grid combines 140 computing centres in 33 countries;[2]
- The Folding@Home volunteer computing project has over 430.000 cores online.[3]

Therefore, a fully centralised approach to resource allocation may not be feasible, as the central resource broker will become a bottleneck for system performance and presents a single point of failure (Wolski et al., 2003). Furthermore, such settings are not necessarily cooperative: stakeholders may have conflicting interests and may be motivated by their individual profit (e.g. stakeholder A and stakeholder B each have their own computational workflow that they want to be completed as soon as possible, hence they are in conflict over who gets to use the available resources; a resource provider may sell the use of its data center for a profit). Therefore, an approach that acknowledges the autonomy of the different actors within a Multi-Agent System (MAS) is required.

In more detail, if we consider a truly open infrastructure, there may be a very large number of agents providing a certain resource and a large number of agents that need such a resource. For a number of reasons, some agents may

[1] http://www.top500.org/
[2] http://lcg.web.cern.ch/LCG/
[3] http://fah-web.stanford.edu/cgi-bin/main.py?qtype=osstats

be more reliable (i.e. more likely to provide full use of the resource, or to settle the payment) than others. We refer to this problem as *execution uncertainty*. For example, a desktop computer providing its idle CPU time will typically be less reliable than a dedicated machine in a data center with equivalent CPU power, because a computation on a desktop machine may be halted when the owner resumes use of the machine, while a resource provider that specialises in selling its computational power will strive to prevent unnecessary interruptions. The situation is complicated further by the fact that agents may enter or leave the system at any time, which means that for an allocation to be meaningful, it will need to be completed in a small time frame, since it may be invalidated at any time by some of the participating agents leaving the system. Similarly, given that agents may have strict deadlines regarding task execution, a timely allocation is important.

Now, in a setting where autonomous agents compete for a limited demand and supply of a resource, market-based resource allocation mechanisms are a natural choice, since they are designed such that desirable overall system behaviour emerges from the agents' selfish, profit-motivated behaviours (Clearwater, 1995). In addition, as mentioned above, we want to avoid the need for a centralised resource broker. Hence, we need a decentralised, market-based resource allocation mechanism. Considering these requirements, the Continuous Double Auction (CDA) is an appropriate choice (Dash et al., 2007).

However, the CDA is not designed with *execution uncertainty* in mind. When there is uncertainty about the reliability of interacting agents, an agent A can represent its beliefs about the reliability of agent B by its *trust* in agent B. Therefore, in this thesis we propose an extension of the CDA that allows agents to use trust information in their decision making.

Given this, in Section 1.1, the CDA is introduced in detail. Section 1.2 explains how the CDA fails in a setting where execution uncertainty is present and argues that a solution based on *trust* is required. Then, the research goals of this work are detailed in Section 1.3 and the research contributions are summarised in Section 1.4. Finally, Section 1.5 provides an overview of the structure of this thesis.

1.1 The Continuous Double Auction

In the CDA, both buyers and sellers may submit their bids (offers to buy) and asks (offers to sell) to the market at any time during the trading period. The market clears *continuously*, that is, whenever a transaction is possible. In the single-unit CDA, this is whenever the highest bid is at least as high as the lowest ask. The CDA is a *double* auction, because there can be both multiple buyers and multiple sellers. Typically all messages sent by agents in the market are made public, anonymously (i.e. without making the identity of the sender public).

Specifically, in the CDA, the market simply collects and emanates information from and to traders. It maintains ordered lists of the current bids and asks: the *order books*. The clearing process can be implemented by simply checking whether the most recently submitted bid is at least as high as the current lowest ask, or the most recently submitted ask is at least as low as the highest bid. If that is the case, the market clears, i.e. a transaction takes place (see Section 2.2

for a more in-depth discussion of the CDA).

This contrasts with centralised mechanisms, that will typically collect bids and asks over a certain fixed period of time and then calculate an allocation as one big optimisation problem. In the CDA, on the other hand, the allocation of resources emerges from the interactions between traders in the market. As a consequence, the CDA offers very little in the sense of guarantees regarding the optimality of the allocation and there is no known optimal trader behaviour. Therefore, it is up to traders to adopt a strategy (i.e. a systematic way of making trading decisions) that ensures they get a good payoff from their participation in the mechanism.

Notably, the Zero-Intelligence (ZI) strategy was developed by Gode and Sunder (1993) to show that in the CDA, even agents that shout random prices (and thus have no intelligence) achieve market efficiency that is close to that of human traders. Hence, the CDA is inherently efficient. Subsequently, Cliff and Bruten (1997) argued that Gode and Sunder's results depend on a specific market structure, and that to achieve human-like performance in general, a more intelligent strategy is needed. To this end, they developed Zero-Intelligence Plus (ZIP), a minimally intelligent adaptive strategy for the CDA. Many more strategies have subsequently been developed that attempt to achieve even higher market efficiency. Trading strategies are discussed in more detail in Section 2.3.

Specifically, market efficiency is defined in terms of the *social welfare* achieved by the system. Social welfare is defined as the sum of the utility (or profit) derived by each of the individuals in the system. Then, market efficiency is defined as the ratio of the social welfare achieved by the market to the social welfare of the optimal allocation.

The computational complexity of the CDA, from the point of view of the market, is very low. In fact, the CDA can even be implemented in a fully decentralised fashion, not requiring a central market or auctioneer at all. Specifically, Ogston and Vassiliadis (2002) compare the scalability of the CDA with a central auctioneer, to a CDA with a distributed hierarchical set of auctioneers to a fully decentralised CDA implemented as a peer-to-peer system. They show that the peer-to-peer system performs better than even the hierarchical set of auctioneers for systems with more than 5,000 traders, with no considerable loss in market efficiency. Moreover, the peer-to-peer implementation scales beyond 160,000 traders, as the communication cost remains constant, while it is linear for the central auctioneer. Thus, the CDA may be implemented as a more or a less centralised system based on the needs of the specific application and can scale to an arbitrary number of traders.

In summary, the following properties make the CDA appropriate for our setting (large scale, open, distributed, computational resource allocation):

- The CDA can allocate among multiple buyers and sellers;
- Because the market clears continuously, the CDA can inherently deal with agents entering and leaving the market during the trading process and with new demand and supply continuously appearing;[4]
- The CDA can scale to very large numbers of buyers and sellers;

[4]That is not to say that the trading process will be unaffected by these events – indeed later we will see that a good trading strategy for the CDA will need to be able to quickly adjust to such changes.

- Because the demands on the central entity can be made as low as is required by the application, i.e. the mechanism can be fully decentralised, the CDA can avoid the dangers of a single point of failure and being a bottleneck for system performance inherent in centralised mechanisms.

Thus, the CDA exhibits highly desirable features for our domain. However, the CDA in its standard form is not robust to execution uncertainty, as is explained in the following section.

1.2 Execution Uncertainty

Now, we consider one aspect of computational resource allocation in a truly open infrastructure, with which the CDA cannot deal: *execution uncertainty*. That is, successful *execution* of an agreed transaction cannot be guaranteed in this setting. Specifically, an agent offering a certain resource may fail to provide access to that resource (e.g. due to systems failure, or interruption to perform a task of higher priority). Similarly, an agent that has agreed to pay a certain amount may fail to complete the payment. Typically, different agents will have varying degrees of *reliability*, the likelihood with which the agent will provide full use of the resource promised, or the likelihood with which the agent will settle the payment.

Given execution uncertainty, agents trading in the CDA will make sub-optimal decisions, because the CDA does not provide the means to differentiate between transaction partners based on their identity. Trade is conducted purely on the basis of price. With execution uncertainty, this means that a buyer will always choose a low-priced offer that is almost certainly faulty, over a reliable offer that is priced slightly higher. Now, the elicitation, representation and use of such reliability information is covered by models of trust (Ramchurn et al., 2004). Specifically, a trust model allows an agent to gather and represent information on the reliability of other agents in a systematic way. Hence, using a trust model will allow an agent to take (its best estimate of) the reliability of a potential transaction partner into account. Therefore, we believe it is useful and interesting to see whether models of trust can positively contribute to a trading agent's success (profit or utility) in a continuous trading environment, by allowing trading agents to balance cost and reliability of the transactions they agree to. However, in the traditional CDA, agents cannot use a trust model, because they do not control the clearing process by any other means than the prices they shout. Moreover, in the CDA, bids and asks are anonymous and thus a trust model cannot be used to judge the value of a shout, since trust is based on the identity of the buyer or seller. Thus, a new mechanism is required.

To this end, in this thesis, we propose a novel variant of the CDA, the Trust-Based CDA (T-CDA), that allows agents to use a trust model in their decision making process to assess whether to accept or reject offers based on cost and the reliability of the proposer. Additionally, we develop Trust-Based ZIP (T-ZIP), a rudimentary trading strategy for the T-CDA. In so doing, we hypothesise that this mechanism with this strategy will be robust to execution uncertainty and will allocate resources in an efficient manner.

1.3 Research Objectives

The aim of this thesis may be summarized in the following research question:

> May the Continuous Double Auction be extended by a component that enables agents to balance cost and reliability of the transactions they agree to, by incorporating a trust-model in their decision making, in a way that achieves close to optimal social welfare even when faced with execution uncertainty?

Against this background, the following may be identified as the four main research objectives for this thesis:

1. To create a novel trading mechanism, based on the CDA, that is robust to execution uncertainty, by allowing agents to use a trust model in their decision making during the trading process.
2. To study the properties of this new mechanism with minimally intelligent traders. This includes efficiency (social welfare in comparison to the optimum social welfare), individual rationality of participating in the mechanism, balance of utilities derived by buyers and sellers, and robustness against unreliable trust models.
3. To implement a trading strategy for the new mechanism.
4. To study the properties of the developed trading strategy and to benchmark its performance.

Each of these objectives is addressed by this thesis, with the overall aim of developing a decentralised resource allocation mechanism that is robust to execution uncertainty.

1.4 Research Contributions

Given the research objectives outlined above, they are addressed in this thesis through the following contributions:

1. *The Trust-Based CDA (T-CDA) mechanism.* A novel trading mechanism, based on the CDA, that allows traders to use their trust model in making trading decisions. The mechanism is empirically investigated using the ZI strategy. The T-CDA is shown to be robust to execution uncertainty, if traders are given perfect and complete trust information, whereas the CDA is shown to break down, because it does not allow agents to use trust information. Furthermore, the robustness of the mechanism to unreliable trust information is empirically demonstrated.
2. *The Trust-Based ZIP (T-ZIP) strategy.* A rudimentary strategy for the T-CDA that is based on the ZIP strategy. The T-ZIP strategy is not a generally applicable trading strategy, but rather is used to further investigate the T-CDA mechanism and to more clearly identify the requirements for a completely general trading strategy. The properties of the T-ZIP are empirically investigated and benchmarked against the ZI strategy.

1.5 Thesis Structure

Next, an overview of literature on market-based resource allocation, trust in market-based systems and on the CDA is provided, in Chapter 2. Against this background, in Chapter 3, the problem is formalised and some of its properties are analysed. Furthermore, suitable measures for empirical evaluation are identified and the research questions are framed within the problem model.

Chapter 4 describes the design and motivation of T-CDA mechanism and its implementation in a simulated environment. An empirical evaluation is performed that shows the new mechanism is an improvement over the traditional CDA. In Chapter 5, a rudimentary trading strategy is developed, based on the ZIP strategy. It is empirically compared to the ZI behaviour and to the ZIP in the traditional CDA. Finally, Chapter 6 concludes and identifies directions for further work.

Chapter 2

Literature Review

In this chapter, an overview of relevant previous work is given. First, in Section 2.1, open distributed systems and specifically the Grid are briefly introduced, and it is shown that the direction taken by this thesis fits well with current work on Grid resource allocation and indeed is a useful addition to it.

Then, a detailed description of the CDA is given and some important previous work is summarised in Section 2.2. Section 2.3 explores the relevant work on trading strategies for the CDA. Finally, Section 2.4 summarises the material discussed in this chapter.

2.1 Background

As was noted in Chapter 1, the motivation for this thesis comes from resource allocation for large-scale, open, distributed systems and, specifically, Grid computing. Therefore, this section provides some additional background on the Grid and the precedents for market-based resource allocation in the Grid. Then, a brief overview of auctions, as used for resource allocation, is provided. This is followed by an introduction on trust in multi-agent systems and previous work incorporating trust in market-based mechanisms. Finally, an example scenario is discussed that further motivates the need for trust in resource allocation for open distributed systems.

2.1.1 The Grid

Research on the *Grid* (Foster and Kesselman, 2003) aims to make computation at the enormous scale required by modern science (and enterprise) possible. For example, no single institution has the computational, storage or manpower capacity to store and analyse the amount of data that is produced by experiments carried out by the Large Hadron Collider. Therefore, the institutions involved in the LHC experiments need to share their resources and coordinate their problem solving. Grids provide an infrastructure to discover, combine and use resources regardless of the details of the underlying hardware or their geographical location. In Grid terminology, resource sharing is done in the context of a *Virtual Organisation* (VO), in which several real-world organisations may come together to share their resources towards a common end, subject to certain terms or conditions.

In other words, access to computational resources is remodelled according to a *utility computing* paradigm analogous to the electricity grid, where computational power becomes pervasive and available on-demand (Foster and Kesselman, 2003). Essentially, software design is decoupled from the underlying hardware, its geographic location and its ownership.

Grid research has focussed on producing "specifications and technologies realising service-oriented architectures according to robust distributed system principles" (Foster et al., 2004). There has been less emphasis on mechanisms that deal reliably with failure and that can adapt to changing conditions. This is the case because up to this point, Grid technology has been used mainly in cooperative settings (Chevaleyre et al., 2006), where several organisations work together towards a common end.

However, as Grid applications become more wide-spread and the distributed systems they build become more open, there is a need for more flexible, autonomous reasoning entities that use intelligent problem solving to achieve their goals, i.e. agents (Foster et al., 2004). As the number and variety of participants in Grid systems increase, so does the potential for conflicting interests. Therefore, the cooperative model of sharing of resources, used by Grid solutions thus far, becomes less appropriate and a competitive model is desirable, i.e. economics (Wolski et al., 2003).

Fortunately, within the multi-agent systems (MAS) research community, significant work has already been done to bring concepts from economics and MAS together. In particular, game theory, a branch of micro-economics, has long been a tool in MAS (Wooldridge, 2002). More recently there has been a move towards computational mechanism design (Dash et al., 2003), which integrates ideas from game theory and distributed systems theory to provide a foundation for the design of real-world, tractable, MAS.

Specifically, computational mechanism design has been applied to resource allocation, creating market-based mechanisms that allocate resources between noncooperative agents. Such systems allow for resource allocation in a context where agents have conflicting interests. The challenge for the mechanism designer is to achieve good system-wide properties despite the fact that agents act selfishly. Indeed, market-based resource allocation has been applied to the Grid (Buyya et al., 2000, 2005, Gomoluch and Schroeder, 2003, Wolski et al., 2003, 2001). A comprehensive overview of market-based resource allocation in computational Grids has been provided by Buyya and Bubendorfer (2009).

As can be seen, much work has gone before that applies market-based techniques to computational resource allocation and the Grid. However, unlike the work in this thesis, execution uncertainty is not addressed by the resource allocation mechanisms in that work.

2.1.2 Auctions

Quite often, market-based mechanisms take the form of auctions. There are two basic types of auctions considered. The first is the *clearing house* type of auction (Krishna, 2002), where a central auctioneer gathers all bids and does a 'one shot' calculation to determine an allocation. The prototypical example of such a mechanism is the Vickrey-Clarke-Groves (VCG) mechanism. This mechanism has been extended in several ways to make it more suitable to certain types of allocation scenario (Dash et al., 2007, Porter et al., 2008), and has been applied to Grid resource allocation (e.g. Schnizler et al., 2008). The advantage of this mechanism is that the best strategy for agents is to bid their true valuation and that it finds the best (most efficient) allocation possible, under the usual assumptions of game theory.

The second type of auction mechanism is more decentralised and trade goes on continuously as, for example, on the stock exchange. The allocation is not calculated in 'one shot' by the auctioneer, but rather is determined by the market dynamics. Here, the prototypical example is the Continuous Double Auction (Smith, 1962). In this type of auction, the guarantees of the VCG do not hold. So agents may adopt a strategy and the most efficient allocation is no longer guaranteed to be found. However, even if agents adopt the very simple Zero Intelligence (ZI) strategy, the market still finds relatively efficient allocations (Gode and Sunder, 1993). The CDA, like the VCG, has been extended to fit different usage scenarios (e.g. Dash et al., 2007) and has been applied to Grid resource allocation (Buyya et al., 2005, Pourebrahimi et al., 2006, Tan and Gurd, 2007).

Thus, there is notable previous work on adaptation of auctions and, specifically, the CDA to specific circumstances and to Grid resource allocation in particular. Hence, our proposal fits within this tradition.

2.1.3 Trust

We define trust as the estimate one agent has about the reliability of another. The trust an agent places in others may be modelled in several competing ways. One such way is grounded in probability theory, often using some form of Bayesian inference (Ramchurn et al., 2004). These trust models can estimate the probability of different outcomes of a transaction with a certain agent. In a market setting, one of the advantages of using probabilistic methods is that the estimated probabilities of different outcomes may be used to calculate the expected utility of the outcome. Hence, they integrate readily with the decision-theoretic means of making decisions: choose the action that leads to the highest expected utility (reward, value or profit). For this reason, we assume that trust is modelled in a probabilistic fashion.

Now, because in a Continuous Double Auction, trade is continuously being conducted, there is an opportunity for agents to learn from each others' actions in the market. Specifically, they may learn about the reliability of other agents. In more detail, such trust may be built in three main ways (Ramchurn et al., 2004):

- Learning: The agent learns how reliable each other agent is through direct interaction with them;
- Reputation: The agent asks other agents in its environment to provide an estimate of reliability of each other agent;
- Socio-cognitive: The agent bases its trust estimate on several estimates of socio-cognitive properties[1] of the other agent.

[1]Socio-cognitive models adopt a higher-level view of trust that takes the knowledge of motivations of other agents for granted and proposes ways to reason about these motivations.

The work in this thesis does not depend on any specific way of acquiring trust. However, when applied to large-scale systems, agents will need to use more than just learning, since having direct interactions with a large proportion of the population is not feasible.

In addition to having several ways of building trust, different matters complicate this picture. For example, in order to learn to trust an agent, we must be certain that the agent we are interacting with is truly the agent it says it is. Hence, agents must be authenticated. In the context of Grid computing, we may take this problem to be solved (Foster et al., 1998). In the case of reputation, we must determine (i) how to gather ratings from other agents, (ii) how to aggregate these ratings into knowledge about trustworthiness and (iii) how to ensure that ratings are provided truthfully (Ramchurn et al., 2004). Even when ratings are provided truthfully, aggregating them can be troublesome; one may have to deal with the absence of information about an agent, or with the fact that agents may rate each other differently because they have different preferences. Issue (iii) must be dealt with at the system level: the mechanism that is used to elicit ratings should ensure that agents that provide untruthful reports are punished (receive diminished utility).

However, it must also be noted that market interactions are not necessarily the only or the primary source of trust information for agents. Reliability information may be provided by an external source (e.g. an independent company that surveys different providers). The above discussion is intended to make clear that trust information acquisition is a difficult matter and indeed that we should not take the availability of *accurate* reliability estimates for granted. Therefore, it is important for a trust-based market mechanism to be robust to inaccuracies in the trust information.

Returning to market-based mechanism and auctions, it must be noted that for centralised auctions, there has already been some degree of success in the integration of trust in specific mechanisms (Dash et al., 2004, Porter et al., 2008). However, to date there is no work on integrating trust into decentralised mechanisms such as the CDA.

2.1.4 A Motivating Scenario

Today, most Grid implementations, although more open and cross-institutional than traditional distributed systems, still work in a fairly constrained, benevolent, environment. However, as Grid technologies become more accessible and wide-spread, this may change. Companies may be interested in selling their data center overcapacity during off-hours to third parties. Research institutions may be interested in buying extra capacity to shorten simulation times. In this way, a competitive market may be set up, that contrasts with the currently predominant cooperative approach. This market may either trade resources for 'real' money, or may induce an artificial economic system of its own (Wolski et al., 2003).

Projects like SETI@home (Anderson et al., 2002) have established that consumer computers can be a valuable computational resource for academics and that consumers are willing to allow the use of that resource, even without direct monetary compensation. As the Grid matures, there is no reason why consumers would not enter the computational utility market as producers. They would provide small amounts of resource per consumer and very variable reliability, somewhat analogous to wind mills in the electricity grid. Consumers of computational resources would have to be able to find these producers in large numbers in order to farm out their computational needs. The decrease in reliability may be compensated for by decreased unit cost, allowing strategies such as those proposed by Stein et al. (2008), provisioning critical parts of a work flow to several providers.

In the context of Grid systems, it is natural to imagine brokers that would 'virtualise' over these large numbers of producers, by simply agreeing to complete some task by a certain deadline and provisioning most of the workload to consumer computers, possibly complementing this with data center capacity where greater reliability is needed, for example when a task is critical for meeting the deadline and time is running short. The point is, that it is quite possible that excess consumer computational power can be used in the Grid with the same ease as the overcapacity of a large data center, by allowing intermediaries to boost reliability and eliminate the additional difficulty of finding and provisioning resources in these potentially very small amounts. As per the spirit of the Grid, the consumer computer pool may be virtualised into a larger pool.

Given this scenario, a market mechanism is needed that is robust to the appearance and disappearance of traders at any moment, that responds well to changing market conditions, that allows very large numbers of traders to interact and that is reasonably efficient. Furthermore, it should be attractive to both buyers and sellers of resources to enter the market. For markets with no execution uncertainty, the CDA is such a mechanism. However, the CDA does not allow agents to balance costs with the risk of failure. Using the CDA, a resource broker as considered above, would either need to consistently bid very low prices (risking not acquiring a resource at all), or run the risk of paying too much for an unreliable resource. Therefore, in order for this scenario to be realisable, a mechanism is needed that allows agents to take their trust in others' reliability into account.

If a truly open Grid infrastructure is to be realised, we can assume virtually no control over individual providers and consumers and hence there are very few means of ensuring that individuals act reliably. Therefore, the market dynamics themselves must ensure that reliability is rewarded. Allowing agents themselves to choose their trading partners based on their perceived reliability is a way to achieve this.

In this section, the relevant background on the Grid and open distributed resource allocation were provided, as well as an overview of auctions for resource allocation and an introduction on trust models. Moreover, a motivating scenario was discussed, in which resource allocation must be decentralised and competitive, while still being efficient, making the CDA an appropriate choice. However, the mechanism must also be robust to execution uncertainty. Therefore, the mechanism proposed in this thesis extends the CDA so that it is robust to execution uncertainty. The following two sections provide the necessary background on the CDA.

2.2 The Continuous Double Auction

Here we provide a brief overview on the micro-economics of markets (adapted from Vytelingum, 2006, chap. 2), followed by a description of the CDA mechanism. The next section will go into agent behaviours within the CDA, the *trading strategies*.

In a market, demand is defined as the willingness or ability of a consumer to purchase a given resource. The demand curve represents the amount of a resource that buyers are willing and able to purchase at various prices. Conversely, supply is the willingness or ability of a producer to provide a given resource. The supply curve represents the amount of resource that producers are willing and able to provide at various prices. In Figure 2.1, demand and supply curves for a specific market are superimposed. The demand and supply curves meet at the *competitive market equilibrium*:

Definition 2.1 (Competitive Market Equilibrium)**.** This is where demand meets supply in a free market populated by profit-motivated selfish agents. The *competitive equilibrium price* is the corresponding price q^*. The transaction prices in the CDA are expected to converge towards q^*. The equilibrium is *competitive* because it is the competition among buyers and sellers that drives transaction prices to q^*. The corresponding quantity v^* is the *market equilibrium quantity.*

At the competitive market equilibrium price q^*, the social welfare (defined in Section 1.1) of the system is maximised. Now, in Figure 2.1, because the demand and supply curves intersect over a range of quantities, we have a *volume tunnel*, where the equilibrium quantity can be $v^* - 1$ or v^*. However, we will assume that goods are desirable and thus, the equilibrium quantity is v^*. In Figure 2.2, the demand and supply curves intersect over a range of prices and hence there is a *price tunnel* between q^*_s and q^*_b. q^* lies somewhere within this range.

We call the price at which each agent is willing to transact its *limit price*:

Definition 2.2 (Limit Price)**.** The maximum bid a buyer is currently willing to offer, or the minimum ask a seller is willing to offer.

- ℓ^b_i is the limit price of buyer i;
- ℓ^s_j is the limit price of seller j – may also be referred to as *cost price.*

Now, if we have *complete* and *perfect* information of the market demand and supply, we can maximise social welfare by determining which agents will trade at what price. Agents on the left of the equilibrium point (buyers with limit prices $\geq q^*$ and sellers with limit prices $\leq q^*$) are known as *intra-marginal traders* and will be trading at price q^*. Agents with limit prices equal to q^* will be trading at zero profit, as we assume that goods are desirable. The other traders are called *extra-marginal*, because their limit prices are too high (sellers) or too low (buyers) to trade in the market. Given this optimal allocation of resources, we can define the *efficiency* of other allocations:

Definition 2.3 (Market Efficiency)**.** The ratio of the sum of all agents' utilities in the market to the maximum possible sum of utilities that would be obtained given the optimal allocation.

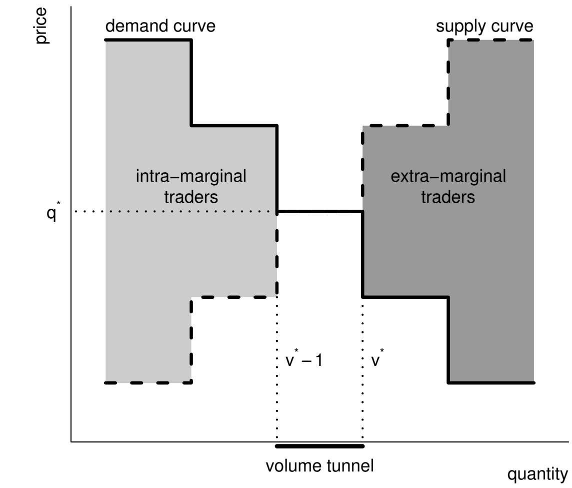

Figure 2.1: Demand and supply curve. Reproduced from Vytelingum (2006).

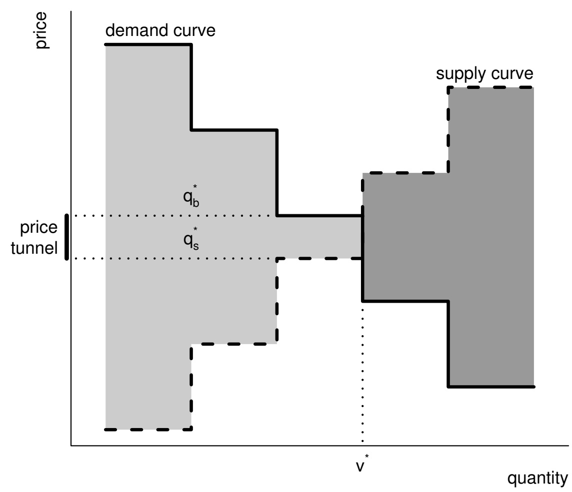

Figure 2.2: Demand and supply curve. q^* lies between q_s^* and q_b^*. Reproduced from Vytelingum (2006).

Smith (1962) demonstrates that markets governed by the CDA mechanism and populated by selfish and profit-motivated (human) traders, can achieve close to optimal market efficiency. Moreover, there is an equilibration of transaction prices to the competitive equilibrium price q^*. It is also demonstrated that if there is a *market shock* (a sudden change in demand and supply at the beginning of a trading day), transaction prices would converge to the new competitive equilibrium price. Convergence of transaction prices to the equilibrium price is measured by the *coefficient of convergence*, α:

$$\alpha = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\tilde{q}_i - q^*)^2}}{q^*} \tag{2.1}$$

where the $\tilde{q}_i$ give a history of n transaction prices. α may also be considered as a measure of *price volatility* in the market.

Thus far, the micro-economics of demand and supply have been discussed and we saw how the optimal (social welfare maximising) solution to an allocation problem (expressed as demand and supply) can be found. Now, we describe the CDA in detail.

In a CDA market, trade is conducted during a *trading day*, the period between opening and closing of the market. In Smith's model, at the beginning of a trading day, traders are endowed with a set of goods to buy or sell, which determines demand and supply. Buyers and sellers submit their bids and asks, respectively. Collectively, these are called *shouts*. If a shout conforms to the *shout accepting rule*, it is placed in the relevant *order book*. Then, the *clearing rule* determines whether a transaction takes place, at a price determined by the *pricing rule*. Whenever a change takes place in the market, the *information revelation rule* determines what information is made public. The rules are defined in detail in Box 2.1.

The CDA may be seen as consisting of two components. First, the *bidding component*, manages the agents' interaction with the order books, through the shout accepting rule. Second, the *clearing component* determines how transactions arise, through the clearing and pricing rules. This is visualised in Figure 2.3.

2.3 Trading Strategies

Having introduced the CDA in some detail, we now turn to the agent behaviour. This is captured by an agent's trading strategy, i.e. the systematic way of making trading decisions in the market place that the agent adopts. A wide variety of trading strategies have been developed over the years. The focus here is on two important strategies, the Zero-Intelligence and the Zero-Intelligence Plus strategies, because of their emphasis on developing a *minimally intelligent* strategy (hence their names) that achieve desirable behaviour in the CDA. Hence, they are important in evaluating the effectiveness of the mechanism. After a thorough description of these two strategies, some interesting other strategies are briefly reviewed, to give an idea of the wide range of possible CDA strategies and of the research that is being conducted into the CDA mechanism.

The market protocol that defines the CDA consists of a number of simple rules. In order to keep track of the offers that have been made, bids and asks are queued into *order books*, which are sorted lists of orders. Bids are sorted from highest to lowest, asks from lowest to highest. The following rules define the CDA protocol in detail:

Shout Accepting Rule Determines which bids and asks are allowed in the market. Primarily, the price must be within the interval $[0, q_{\max}]$. $q_{\max}$ is the maximum bid or ask allowed in the market, to prevent unreasonably high asks and speed up the trading process.

Furthermore, the commonly implemented NYSE shout accepting rule imposes that a new shout must improve upon the current best shout by that agent. When a trader submits a new shout, provided that it improves upon the current shout by that trader, the current shout is simply replaced by the new one.

Information Revelation Rule Determines what information is published to buyers and sellers. Typically, this is current bid and ask prices.

Clearing Rule The market *clears* continuously, whenever the highest bid price is at least as high as the lowest ask. Then a transaction takes place, at a *transaction price*, determined according to the pricing rule. The matched shouts are removed from the order books.

Pricing Rule Determines the transaction price. The average of the matched bid and matched ask prices is typically used in the CDA.

Box 2.1: The CDA protocol

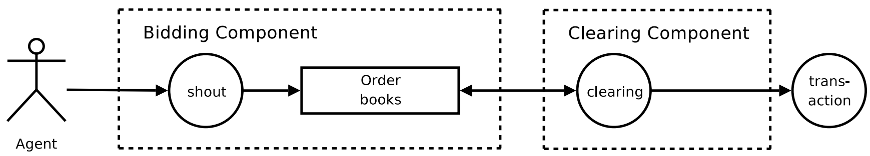

Figure 2.3: The traditional CDA can be seen as consisting of two components. This figure visualises how information flows through these components. Circles show concrete pieces of information in the system. Note that a ‘clearing’ here is a matching of a bid and an ask (not to be confused with the *clearing rule*). The rules govern the information flow visualised by the arrows in this figure. For example, the flow of a shout into the order book (second arrow) is governed by the *shout accepting rule*.

2.3.1 Zero-Intelligence

The Zero-Intelligence (ZI) strategy (Gode and Sunder, 1993), introduced in Section 1.1, is the baseline strategy for the CDA. A ZI agent is not motivated by profit and ignores all market conditions when submitting a bid or an ask. Rather, it will draw a shout price from a uniform distribution with a given range. Gode and Sunder (1993) consider two types of ZI agents: the *unconstrained* Zero-Intelligence Unconstrained (ZI-U) agents, for whom the price range is $[0, q_{\max}]$ and the *constrained* Zero-Intelligence Constrained (ZI-C)[2] agents, which are not allowed to trade at a loss. Therefore, the range for a ZI-C buyer i is $[0, \ell_i^b]$ and for a seller j it is $[\ell_j^s, q_{\max}]$.

Gode and Sunder (1993) show that ZI-C agents exhibit behaviour that is much more like that of human traders than ZI-U agents do. With ZI-C, there is a slow convergence of transaction prices to the theoretical equilibrium and market efficiency is very close to that achieved by human traders. Given this, it appears that market efficiency is almost entirely a result of market structure. Therefore, the previous assumptions that the efficiency of human markets is a consequence of human intelligence (Smith, 1962) is called into question.

However, human traders do have the lowest profit dispersion (i.e. the least variation in individual profits) when compared to the ZI-C and ZI-U traders. From this, Gode and Sunder (1993) note that individual aspects of market performance may be more sensitive to human intelligence than market efficiency.

2.3.2 Zero-Intelligence Plus

The ZIP strategy (also discussed in Section 1.1) is based on the idea that any offer being made and every transaction occurring is an opportunity for an agent to learn how to calibrate its own pricing. To this end, in addition to the *limit price* ℓ_i, agents set a *profit margin* μ_i. Together these determine the shout price q_i:

$$q_i = \ell_i(1 + \mu_i) \ . \tag{2.2}$$

This means that a seller's margin is raised by increasing μ_i and lowered by decreasing μ_i and that $\mu_i \in [0, \infty)$. Buyers raise their margin by decreasing μ_i and lower their margin by increasing μ_i, with $\mu_i \in [-1, 0]$.

The agents must learn the appropriate profit margin from market events. This raises two questions: first, when is it appropriate to raise or lower the profit margin? Second, how should the profit margin be updated? The first issue is addressed by the bargaining mechanism, and the second by the adaptation mechanism.

Bargaining Mechanism

When considering whether to raise or lower its profit margin, an agent has four factors to consider. First, whether it is *active* in the market. This is the case when it is still capable of making a transaction, otherwise it is *inactive*. The other three factors are properties of the last shout: its price q, whether it was a bid or an ask and whether it resulted in a transaction or not. Furthermore, let

[2] In the remainder of this thesis, when we refer to the ZI strategy, this may be taken to mean the ZI-C strategy.

q_i represent the price that agent a_i intended to shout, not taking into account the information from the current last shout.

Whenever a shout is submitted in the market, a ZIP trader will evaluate its bargaining rules (given in Algorithm 2.1 and Algorithm 2.2) to decide whether its profit margin should be updated. If the margin is to be updated, the adaptation mechanism is invoked.

Algorithm 2.1 Bargaining algorithm for seller s_i

```
if Last shout resulted in a transaction at price q then
  if q_i ≤ q then
    raise profit margin
  end if
  if Last shout was a bid AND s_i is active AND q_i ≥ q then
    lower profit margin
  end if
else
  if Last shout was an ask AND s_i is active AND q_i ≥ q then
    lower profit margin
  end if
end if
```

Algorithm 2.2 Bargaining algorithm for buyer b_i

```
if Last shout resulted in a transaction at price q then
  if q_i ≥ q then
    raise profit margin
  end if
  if Last shout was an ask AND b_i is active AND q_i ≤ q then
    lower profit margin
  end if
else
  if Last shout was a bid AND b_i is active AND q_i ≤ q then
    lower profit margin
  end if
end if
```

Adaptation Mechanism

The profit margin of agent a_i, μ_i, is updated according to a delta rule. This is a learning rule that gradually adapts the variable to be learned towards its desired value based on the inputs it receives. Let $\mu_i(t)$ be agent a_i's profit margin at time t and $q_i(t)$ its calculated shout price at time t. Then we update the margin μ_i on the transition from time t to $t+1$ as follows:

$$\mu_i(t+1) = \frac{(q_i(t) + \delta_i(t))}{\lambda_i} - 1 \ , \tag{2.3}$$

parameter	range
$\mathcal{R}_i(t)$ (increase)	$[1.0, 1.05]$
$\mathcal{R}_i(t)$ (decrease)	$[0.95, 1.0]$
$\mathcal{A}_i(t)$ (increase)	$[0.0, 0.05]$
$\mathcal{A}_i(t)$ (decrease)	$[-0.05, 0.0]$
β_i	$[0.1, 0.5]$
γ_i	$[0.0, 0.1]$
$\mu_i(0)$ (sellers)	$[0.05, 0.35]$
$\mu_i(0)$ (buyers)	$[-0.35, -0.05]$

Table 2.1: Default ranges of ZIP parameters. Each required value is generated from a uniform distribution over the given range.

where $\delta_i(t)$ is the momentum-based delta value. The momentum-based delta value is defined as follows:

$$\delta_i(t+1) = \gamma_i \delta_i(t) + (1-\gamma_i)\Delta_i(t+1) \tag{2.4}$$

$$\delta_i(0) = 0 \tag{2.5}$$

where $\gamma_i \in [0, 1]$ is the *momentum coefficient*, and $\Delta_i(t)$ is the delta value,[3] calculated using a_i's *learning rate* β_i and a *target price* $\tau_i(t)$:

$$\Delta_i(t) = \beta_i(\tau_i(t) - q_i(t)) \ . \tag{2.6}$$

There are many ways in which the target price $\tau_i(t)$ could be set. For standard ZIP traders, the target price is a stochastic function of the shout price $q(t)$:

$$\tau_i(t) = \mathcal{R}_i(t)q(t) + \mathcal{A}_i(t) \ , \tag{2.7}$$

where $\mathcal{R}_i(t)$ is a randomly generated coefficient that sets the target price relative to the price $q(t)$ of the last shout, and $\mathcal{A}_i(t)$ is a small *absolute* price alteration. When the intention is to increase the dealer's shout price, we set:

$$\mathcal{R}_i > 1.0 \ ; \ \mathcal{A}_i > 0.0 \ . \tag{2.8}$$

When the intention is to decrease the price, we set:

$$0.0 < \mathcal{R}_i < 1.0 \ ; \ \mathcal{A}_i < 0.0 \ . \tag{2.9}$$

$\mathcal{R}_i(t)$ and $\mathcal{A}_i(t)$ are randomly generated in an independent and identical way for each individual agent and time step.

Cliff and Bruten (1997) randomly generate each of the many values defined here from uniform distributions with certain ranges. These ranges are given in Table 2.1.

Results

With the ZIP strategy, transactions converge towards the competitive equilibrium price after a few trading days and remain at that level with low variance. The ZIP strategy was shown to achieve results closer to human performance

[3] Note that when $\gamma_i = 0$, $\delta_i(t) = \Delta_i(t)$.

than ZI-C, even though its parameters were not optimised for the demand and supply of the market (Cliff and Bruten, 1997). Moreover, profit dispersion for ZIP traders is much lower than in a market with ZI-C traders. The ZIP strategy is also able to converge to a new competitive equilibrium after a market shock.

2.3.3 Other Strategies

The two strategies described above, ZI and ZIP, are the most important for this thesis, because ZI is used as a lower bound to the performance that can be expected from the mechanism and ZIP forms the basis of the T-ZIP strategy developed in this thesis. Here, to illustrate work that has been done on the CDA and to show that the strategies described above are just examples of what is possible, a number of important alternative strategies are overviewed (summarised from Vytelingum, 2006, chap. 2).

Kaplan The Kaplan strategy does not adapt to market efficiency or infer the market equilibrium, but attempts to exploit the bidding behaviour of other agents by *sniping* at any profitable deal. It will wait while the other strategies do the negotiating, and then, based on some simple heuristics, take away a good deal at the last moment. The Kaplan strategy does well in a heterogeneous environment, but a market populated with only Kaplan strategies does not perform efficiently because the prices will not be driven towards the equilibrium.

ZIP60 The ZIP strategy described previously was subsequently extended to ZIP60 (Cliff, 2005). The original set of 8 parameters for updating the profit margin was extended to 10 and a different set of parameters was used for each of the 6 different learning rules. The 60 parameters are selected through a genetic algorithm optimisation that minimises price volatility (Equation 2.1). The ZIP60 is thus tailored to a specific market and can accomplish significant improvements over ZIP. However, this tailoring to a market means that demand and supply must be known a priori, which is usually not a realistic assumption.

GD The GD strategy (Gjerstad and Dickhaut, 1998) builds a belief function that indicates whether a particular shout is likely to be accepted in the market from the market history. Given this, the bidding strategy submits the shout that maximises the trader's expected utility, which is the product of its belief function and its utility function. The GD strategy takes information recency into account by limiting the trader's memory length. GD was shown to achieve close to optimal efficiency and rapid convergence of prices to the optimum. In heterogeneous populations, it extracted 1.7% more profit than ZIP.

GDX Building on the GD strategy, GDX additionally takes the time left in the trading day into account. This means the GDX trader is able to wait for more profitable transactions that may appear later in the trading day (Tesauro and Bredin, 2002). Given an opportunity to submit a shout to the market, GDX also estimates the number of bidding opportunities before the market closes. It then uses dynamic programming to calculate the optimal shout, taking this estimate into account. GDX was shown to outperform both ZIP and GD.

FL The FL (fuzzy logic based) strategy (He et al., 2003) employs fuzzy reasoning in order to determine the best bid or ask given the current state of the market, based on market history. It defines a set of possible transaction prices represented by fuzzy numbers and uses heuristic rules to infer the best action. Although the FL strategy does well in heterogeneous populations, performance is poor in homogeneous environments, like the Kaplan strategy.

AA The adaptive aggressiveness (AA) strategy (Vytelingum et al., 2008) adopts both short-term and long-term learning to adapt to market conditions. The short-term learning updates the aggressiveness of the bidding behaviour, where more aggressive means more willing to trade off potential future profits for a better chance of transacting. The long-term learning adapts the way a trader's aggressiveness influences its bidding behaviour. AA was shown to outperform both ZIP and GDX in both homogeneous and heterogeneous markets.

While the strategies reviewed in this section are very interesting, they are also a good deal more complex than the ZI and ZIP strategies (with the possible exception of Kaplan - which, by itself, is not a viable strategy anyway). This means not only that they would be more difficult to implement or adapt to a new mechanism, their more intricate behaviours make analysis of experimental results more difficult. Hence, for the initial evaluation of a new mechanism, it is better to employ a simple strategy such as ZI or ZIP. However, the above shows that there is considerable interest in automated trading strategies for the CDA and that any extension of the CDA has an extensive body of work on trading strategies to draw from. Future work might extend any of these strategies and apply them to the T-CDA. See Section 5.4 for directions for such work.

2.4 Summary

In this chapter, an overview of the current state of market-based resource allocation in Grid computing was provided. A scenario was discussed where a decentralised resource allocation mechanism that deals with both non-cooperative agents and execution uncertainty is not just desirable, but absolutely required. To date, such a mechanism does not exist.

Then, the micro-economics of markets were reviewed and it was noted that supply and demand curves meet at an *equilibrium price* q^*, and that allowing agents to trade at this price as long as this does not yield a loss for them will result in the *optimal allocation* (i.e. the allocation that maximises social welfare). Hence, for a standard market scenario, we may calculate the *market efficiency*, an objective measure of the quality of an allocation.

Against this background, the CDA can be viewed as a set of rules that determine how bids and asks by traders eventually result in transactions. Even with ZI traders (which shout prices randomly), the CDA achieves high market efficiency, which thus seems to be attributable mainly to the structure of the market. Furthermore, work on the ZIP strategy has shown that ZI does not always do as well and instead proposes a simple adaptive strategy.

Finally, a number of different strategies were discussed, showing the wide range of strategies that have been proposed for the CDA, as well as their strengths and weaknesses. This shows that there has been extensive work on automated trading in the CDA and that this is still ongoing. Therefore, any extension of the CDA has an extensive body of work to draw from.

Chapter 3

Problem Definition

Having identified the research objectives in Section 1.3 and the relevant previous work in Chapter 2, the problem is formalised in this chapter. First, the trading environment in which the T-CDA mechanism will be evaluated is defined in Section 3.1. Then, Section 3.2 shows that, in general, there is no market equilibrium in this trading environment. Section 3.3 defines the optimal solution to the allocation problem and how to find it. Against this background, Section 3.4 defines desiderata for the T-CDA and its evaluation. Finally, Section 3.5 provides a summary.

3.1 Modelling the Trading Environment

This section formally introduces the problem setting. That is, a model of the trading environment is defined. In this model, a number of simplifying assumptions are made:

- The set of buyers and sellers is fixed;
- No new demand or supply appears during a trading day – hence, the full demand and supply are known at the start of a trading day. Thus, the allocation for each day can be calculated as a single optimisation problem;
- Failure is binary, that is, either failure or success;
- Each agent has only one order to fill;
- The utility functions are of a specific form – agents are risk-neutral, non-malicious and value monetary gain linearly.

These assumptions allow us to calculate an optimal allocation against which to compare the efficiency of the mechanism (Section 3.3). None of these assumptions, however, are required by the mechanism. They merely provide a simple scenario in which to evaluate the mechanism, without loss of generality.

In what follows, first the part of the model in which a traditional trading mechanism (specifically, the CDA) would be evaluated is introduced. Then, the notion of execution uncertainty and how it impacts on this model is discussed. Finally, because agents do not have perfect and complete information of each other's reliability, agents are given a trust function.

3.1.1 Market Definition

We denote the set of buyers as $b_1, b_2, \ldots, b_n \in B$ and the set of sellers as $s_{n+1}, s_{n+2}, \ldots, s_{n+m} \in S$. Then, the set of agents is denoted as $A = B \cup S$. As a convention, we generally refer to a generic buyer as b_i, a seller as s_j and a generic agent as a_i, when we do not distinguish buyers and sellers.

Every agent participating in the market is given an *endowment*. For a buyer, an endowment is an order to buy a single unit of resource for at most the specified *limit price*, ℓ^b_i (Definition 2.2). For a seller, an endowment is an order to sell a single unit of resource for at least the specified *cost price*, ℓ^s_j.

Given their endowments, buyers place *bids* (offers to buy) and sellers place *asks* (offers to sell) in the market. Collectively, bids and asks are referred to as *shouts*. Based on the submitted bids and asks, the market mechanism determines when a *transaction* takes place between a buyer and a seller. We will denote a transaction at price q between a buyer $b_i \in B$ and seller $s_j \in S$ as $t_{i,j}(q)$. After agreeing on a transaction $t_{i,j}(q)$, the buyer pays the seller and the seller transfers some goods to the buyer. The way the shouts are managed in the market can be regimented by different market rules.

3.1.2 Introducing Execution Uncertainty

The setting described above is the one traditionally considered in market-based mechanisms. Moreover, in this work, we do not assume that successful *execution* of a transaction is guaranteed. Instead, we assume that the execution of a transaction is binary, that is, either failure or success[1]. We denote the outcome for the buyer as $e_b \in \{0, 1\}$ and for the seller as $e_s \in \{0, 1\}$. The probability that a buyer is successful (i.e. $P(e_{b_i} = 1)$) is denoted as $p(b_i)$ and that the seller is successful (i.e. $P(e_{s_j} = 1)$) as $p(s_j)$. For example, after $t_{i,j}(q)$, if $e_b = 1$ and $e_s = 0$, buyer b_i has paid for a service, but s_j did not provide that service. In general, every agent a_i is assigned a certain Probability of Success (POS) $p(a_i) \in [0, 1]$, which indicates the likelihood that an agent will honour its agreement.

Given the execution (e_b, e_s) of a transaction, the agents derive utility as follows:

$$
\begin{aligned}
u^b_i(t_{i,j}(q), e_s) &= \begin{cases} \ell^b_i - q & , e_s = 1 \\ -q & , e_s = 0 \end{cases} , \\
u^s_j(t_{i,j}(q), e_b) &= \begin{cases} q - \ell^s_j & , e_b = 1 \\ -\ell^s_j & , e_b = 0 \end{cases} ,
\end{aligned}
\tag{3.1}
$$

where ℓ^b_i is the limit price of b_i (i.e., the maximum b_i is willing to pay) and ℓ^s_j is the cost price of s_j (i.e. the minimum price at which s_j is willing to sell). These functions follow naturally if we assume that agents are not malicious; i.e. regardless of their own success, they will incur the cost associated with the action they agreed to perform. For example, when a buyer pays q and if he receives the goods (or service), which are worth ℓ^b_i to him, he will derive a utility of $\ell^b_i - q$. Otherwise, his utility is $-q$. Given this definition of utility, the

[1]Failure is binary to simplify our analysis, but this work can easily be generalised to be continuous, to reflect partial success or failure if that is appropriate in a given setting.

expected utility of a transaction is given by:

$$
\begin{aligned}
\hat{u}_i^b(t_{i,j}(q)) &= u_i^b(t_{i,j}(q),1)p(s_j) + u_i^b(t_{i,j}(q),0)(1-p(s_j)) \\
&= \ell_i^b p(s_j) - q \\
\hat{u}_j^s(t_{i,j}(q)) &= u_j^s(t_{i,j}(q),1)p(b_i) + u_j^s(t_{i,j}(q),0)(1-p(b_i)) \\
&= qp(b_i) - \ell_j^s \ .
\end{aligned}
\tag{3.2}
$$

That is, the utility of each outcome multiplied by the probability of that outcome, summed over all possible outcomes (i.e. the normal probabilistic interpretation of expected utility).

Note that our model is equivalent to the setting in which the CDA is normally evaluated, when $p(a_i) = 1; \forall a_i \in A$. In that case, the expected utility functions are simply $u_i^b(t_{i,j}(q),1) = \ell_i^b - q$ and $u_j^s(t_{i,j}(q),1) = q - \ell_j^s$.

3.1.3 Trust

Now, we have defined how a trader should evaluate its expected utility, given perfect and complete information. However, since in general we cannot assume that agents have perfect and complete knowledge of each other's POS, agents hold an estimate of the POS of the other agents. Thus, each agent a_i has a *trust function*:

$$
\mathrm{trust}_i : A \rightarrow [0,1] \ , \tag{3.3}
$$

which represents its best estimate of the probability of success for each other agent. So ideally, $\mathrm{trust}_i(a_j) \approx p(a_j)$. This allows a_i to estimate the expected utility $\hat{u}$ (Equation 3.2) of a transaction:

$$
\begin{aligned}
\tilde{u}_i^b(t_{i,j}(q)) = u_i^b(t_{i,j}(q),1)\mathrm{trust}_i(a_j) + u_i^b(t_{i,j}(q),0)(1-\mathrm{trust}_i(a_j)) \\
\tilde{u}_j^s(t_{i,j}(q)) = u_j^s(t_{i,j}(q),1)\mathrm{trust}_j(a_i) + u_j^s(t_{i,j}(q),0)(1-\mathrm{trust}_j(a_i)) \ .
\end{aligned}
\tag{3.4}
$$

It is rational to agree to a transaction only if the estimated expected utility $\tilde{u}_i(t) \geq 0$. Here we remain agnostic to the origin of this trust function; agents might learn the reliability of others through the observation of market interactions, or they could have some outside source of information.

In summary, a model was defined in which trading agents are endowed not only with their private valuation of a resource, but also with their private POS. The POS determines the likelihood that an agent successfully delivers the resource or completes the payment. Since it is unrealistic to assume agents know each other's POS, we assume each trader will have an estimate of this information, represented by a trust function. The notation introduced is listed in Table 3.1.

3.2 Market Equilibria

As was discussed in Section 2.2, in a market without execution uncertainty, the demand and supply curves meet at the competitive market equilibrium (Definition 2.1). At this equilibrium price, the social welfare of the system is maximised. Now, having defined the expected utility functions for the traders (Equation 3.2), we can ask whether such an equilibrium exists for a market with execution uncertainty.

Symbol	Meaning
B	The set of buyers
S	The set of sellers
A	The set of agents $A = B \cup S$
b_i	A buyer $b_i \in B$
s_j	A seller $s_j \in S$
a_i	An agent $a_i \in A$
ℓ_i^b	Limit price of buyer b_i
ℓ_j^s	Limit price (or: cost price) of seller s_j
$t_{i,j}(q)$	Transaction between b_i and s_j at price q
e_{b_i}	Outcome of an execution for b_i
e_{s_j}	Outcome of an execution for s_j
$p(a_i)$	Probability of success of agent a_i (i.e. $P(e_{a_i} = 1)$)
u_i^b, u_j^s	Utility function of b_i and s_j, respectively
$\hat{u}_i^b$, $\hat{u}_j^s$	Expected utility function of b_i and s_j, respectively
$\text{trust}_i(a_j)$	Trust of agent a_i in a_j
$\tilde{u}_i^b$, $\tilde{u}_j^s$	Estimated expected utility function of b_i and s_j, respectively

Table 3.1: Overview of notation

Theorem 3.1. *In the model defined by Section 3.1, in general, a market equilibrium (Definition 2.1) does not exist.*

Proof. A counter-example will show that a single equilibrium price does not exist. Hence, a market equilibrium as per Definition 2.1 does not exist. Note that for a price to be the equilibrium price, no agent may trade at negative expected utility and social welfare must be maximised. Consider the following market:

$$\begin{array}{llll} B = \{b_1, b_3\} & & S = \{s_2, s_4\} & \\ p(b_1) = 1 & \ell_1^b = 1 & p(s_2) = 1 & \ell_2^s = 1 \\ p(b_3) = 0.5 & \ell_3^b = 2 & p(s_4) = 1 & \ell_2^s = 1 \end{array} \tag{3.5}$$

Now, using Equation 3.2, for each pair of buyer and seller, we can calculate a minimum price at which a transaction is possible, by solving $u_j^s(t_{i,j}(q)) = 0$ for q. The maximum transaction price can be found by solving $u_i^b(t_{i,j}(q)) = 0$ for q. In this instance, the solutions are:

$$\begin{array}{ccc} & s_2 & s_4 \\ b_1 & (1,1) & (1,1) \\ b_3 & (2,2) & (2,2) \end{array} \tag{3.6}$$

where each pair represents $(\min, \max)$. Thus, each pair of buyer and seller can transact. In this case, any potential transaction would have zero expected utility, but assuming (as in Section 2.2) that goods are desirable, two transactions should take place. For example, an optimal solution would be $t_{1,2}(1)$ and $t_{3,4}(2)$. Thus, there is no single equilibrium price.

One might object that the above counter-example is defeated if the assumption that goods are desirable is dropped. However, the counter-example holds for any $\ell_2^s = \ell_4^s > 0.5$. To see this, specifically consider $\ell_2^s = \ell_4^s = 0.8$:

$$\begin{array}{ccc} & s_2 & s_4 \\ b_1 & (0.8, 1) & (0.8, 1) \\ b_3 & (1.6, 2) & (1.6, 2) \end{array} \tag{3.7}$$

Again, two transactions should take place. In fact, in this case two transactions are required in order to optimise social welfare. The choice of which pairs of traders transact is arbitrary. Hence, to optimise social welfare we can optimise the expected utility derived from each single transaction, defined as the sum of the individual expected utilities:

$$U_{i,j}(q) = \hat{u}_i^b(t_{i,j}(q)) + \hat{u}_j^s(t_{i,j}(q)) = \ell_i^b p(s_j) - \ell_j^s + q(p(b_i) - 1) \ . \tag{3.8}$$

Note that when $p(b_i) < 1$, we maximise $U_{i,j}(q)$ by minimising q. When $p(b_i) = 1$, q has no impact on $U_{i,j}(q)$. Hence, in the above market, if we choose (b_1, s_2) and (b_3, s_4) to transact, we must choose $t_{3,4}(1.6)$, because $p(b_3) = 0.5 < 1$. We are free to choose $0.8 \leq q \leq 1.0$ in $t_{1,2}(q)$, however. In any case, the transactions take place at different prices. □

The above proof shows that in general, there is no equilibrium price. However, in special conditions an equilibrium could exist. Specifically, if $p(a_i) = 1 \ ; \ \forall a_i \in A$ (i.e. no execution uncertainty), an equilibrium is known to exist. The remainder of this section analyses different cases to give a clear intuition of when equilibria exist.

From the proof, we know that an equilibrium price does not exist if buyer limit prices and POS differ. By a similar counter-example, an equilibrium price does not exist if seller limit prices and POS differ.

Even if we set $p(s_i) = 1 \ ; \ \forall s_i \in S$ and $p(b_i) = p \ ; \ \forall b_i \in B$, for a $0 \leq p < 1$, an equilibrium price does not exist. To show this, let us examine the acceptable prices for a seller s_j:

$$\hat{u}_j^s(t_{i,j}(q)) = qp(b_i) - \ell_j^s \geq 0 \tag{3.9}$$

$$qp - \ell_j^s \geq 0 \tag{3.10}$$

$$q \geq \frac{\ell_j^s}{p} \tag{3.11}$$

Because Equation 3.8 implies that, in order to optimise $U_{i,j}(q)$, we must choose the smallest possible q, a transaction between s_j and b_i should take place at:

$$q = \frac{\ell_j^s}{p} \ , \tag{3.12}$$

which means that the desirable transaction price depends on the seller's limit price and hence that no single equilibrium price exists.

Besides a market without execution uncertainty, we identify one special case in which an equilibrium does exist. Say $p(b_i) = 1 \ ; \ \forall b_i \in B$ and $p(s_i) = p \ ; \ \forall s_i \in S$, with $0 \leq p \leq 1$. Now a buyer b_i would be willing to transact if:

$$\hat{u}_i^b(t_{i,j}(q)) = \ell_i^b p(s_j) - q \geq 0 \tag{3.13}$$

$$\ell_i^b p - q \geq 0 \tag{3.14}$$

$$q \leq \ell_i^b p \tag{3.15}$$

Here, every buyer's maximum transaction price is multiplied by p. However, because buyers have POS 1, we are free to choose any q between the buyer's limit multiplied by q and the seller's limit. Hence, in this case the seller's lower but identical POSs merely causes a shift in the demand curve. Thus, an equilibrium exists.

Undoubtedly, it is possible to construct different constraints on the endowments or POS that force an equilibrium price to exist. However, in the above we identified the major situations based on POS and noted whether an equilibrium exists or not. In general, there is no equilibrium price in the trading environment defined in Section 3.1. Since finding the optimal allocation and, hence, the market efficiency in the CDA depends on finding the equilibrium price (see Section 2.2), a different way of finding the optimal allocation is required.

3.3 Optimal Solution

In this section, we define and find the optimal solution, given complete and perfect information of all agents. This provides an upper bound on the efficiency we can expect from our mechanism. Given our model, we aim to find the allocation that maximises the sum of the expected utilities of the individual agents, subject to certain constraints. Specifically, an *allocation* is a list of transactions that take place between agents. Thus, we need to decide which agents shall transact and at what price. In this section, a linear program that finds the optimal solution is developed.

First, let us consider how to choose the transaction price given that two agents interact. In order to optimise efficiency, we should maximise the sum of the agents' individual utilities (Equation 3.8):

$$U_{i,j}(q) = \ell_i^b p(s_j) - \ell_j^s + q(p(b_i) - 1) \ . \tag{3.16}$$

From the above formula, we see that when the probability of success of the buyer $p(b_i) = 1$, the transaction price q has no influence on the total expected utility of the transaction. However, when $p(b_i) < 1$, a higher transaction price leads to a lower expected utility. Therefore, if we choose q to optimise $U_{i,j}$, sellers will derive negative expected utility. Hence, participation is not individually rational.

To remedy this, we could demand that $\hat{u}_j^s(t_{i,j}(q)) \geq 0$, however when $p(b_i) < 1$, the result will be that sellers will always break even and thus have no incentive to take part in the market. Instead, we demand that the expected utilities of both parties are equal, to achieve a fair distribution of utility between buyers and sellers:

$$\hat{u}_i^b(t_{i,j}(q)) = \hat{u}_j^s(t_{i,j}(q)) \ . \tag{3.17}$$

Substituting Equation 3.2 into Equation 3.17 completely determines the acceptable transaction price:

$$q = \frac{\ell_i^b p(s_j) + \ell_j^s}{1 + p(b_i)} \ . \tag{3.18}$$

This then also determines the transaction utility:

$$U_{i,j} = \ell_i^b p(s_j) - \ell_j^s + \frac{\ell_i^b p(s_j) + \ell_j^s}{1 + p(b_i)}(p(b_i) - 1) \ , \tag{3.19}$$

where the argument q from equation 3.16 is omitted, since there is only one acceptable price q. These $U_{i,j}$ together define a $|B| \times |S|$ matrix U. Now define the decision matrix $T \in \{0,1\}^{|B|\times|S|}$ as follows:

$$T_{i,j} = \begin{cases} 1 & b_i \text{ and } s_j \text{ transact} \\ 0 & \text{otherwise.} \end{cases} \tag{3.20}$$

Our objective is to find the matrix T that maximises the total expected utility in the system, given that every agent transacts at most once. The exact formulation is given in Algorithm 3.1.

Algorithm 3.1 Linear program to find the optimal allocation

Maximize

$$\sum_{i:b_i \in B} \sum_{j:s_j \in S} (U \cdot T)_{i,j} \ ,$$

subject to

$$\sum_{j:s_j \in S} T_{i,j} \leq 1 \ ; \ \forall_{i:b_i \in B} \ ,$$

$$\sum_{i:b_i \in B} T_{i,j} \leq 1 \ ; \ \forall_{j:s_j \in S} \ ,$$

where $\cdot$ denotes the Hademard (entry-wise) product of two matrices.

After translation into a standard notation, this specification can be executed by an integer programming package. Note that some constraints, such as equality of buyer and seller utility (equation 3.17) do not need to be represented explicitly, as they are enforced by the definition of U seen previously. The value being maximized over is called the *objective function*; in this case it is the sum of the expected utilities of all agents. Because buyer and seller utilities are equal, the sum of the expected utilities of all buyers is half that value.

Now we know how to calculate the optimal allocation using perfect and complete information of the playing field, given all of the assumptions of the model discussed in Section 3.1. This gives an upper bound on the performance of our proposed mechanism, which does not depend on the use of perfect and complete information or said assumptions.

In more detail, U gives an upper bound on the performance of the T-CDA, under the constraint that utility is equally distributed between buyers and sellers. However, the solutions the T-CDA finds do not necessarily obey this constraint, since the allocation emerges from the interactions between traders, rather than being determined by a central auctioneer. Therefore, in evaluating the mechanism, we must separately compare both buyer and seller utilities to $0.5U$.

3.4 Desiderata

Given our problem setting, we define a number of *desiderata* that we believe our mechanism should exhibit, based on the research objectives stated in Section 1.3. In particular:

- The market mechanism should be *efficient*: it should maximise the sum of the expected utilities of the individual agents, since we want to maximise social welfare;
- It should also be *individually rational*, i.e. individual agents will not participate in loss-making transactions. This ensures that we do not disincentivise agents from participating in our market;
- Furthermore, an equal and, thus, fair distribution of profits between buyers and sellers is desirable (again to ensure we have approximately equal numbers of each);
- Additionally, since our model incorporates the notion of POS, we desire the mechanism to be robust against agents having an inaccurate representation of each others' POS, since in the real world, it is unrealistic to assume that agents have perfect and complete information about the reliability of other agents.

Given these desiderata, since the addition of execution uncertainty introduces a number of new problems for the mechanism and the traders, empirical evaluation (Section 4.4) should focus on this aspect. Hence, even though we could investigate any number of demand and supply curves, it is more interesting to fix demand and supply and vary the POS we assign to the agents. Moreover, the trust function is a new addition and it is interesting to study how the mechanism responds to different properties of the trust model. Specifically, it should be investigated how performance breaks down as trust information becomes less accurate. This is done in Section 4.4.5.

3.5 Summary

In this chapter, a formal model of the trading environment was developed. Then, it was shown that in general, no equilibrium price exists in this model. Furthermore, the optimal allocation was defined and a method of finding it was given. Against that background, desiderata for the mechanism were made explicit, as well as directions for its evaluation. This further clarifies and makes concrete the research objectives stated in Section 1.3.

Given this, in the remainder of this thesis, the T-CDA mechanism and the T-ZIP strategy are developed and implemented. Both are empirically evaluated within the framework specified by this chapter.

Chapter 4

Trust-Based CDA

The Trust-Based CDA (T-CDA), a new mechanism based on the CDA, is introduced. This mechanism allows traders to take execution uncertainty into account in their decision making, whilst maintaining the decentralised nature of the CDA. First, the new mechanism and the design decisions that were made are outlined in Section 4.1. Then, Section 4.2 defines a baseline trading behaviour, while the way the trading process is simulated is detailed in Section 4.3. The T-CDA simulator is used in Section 4.4 to empirically evaluate the T-CDA mechanism. Finally, Section 4.5 concludes and summarises the main points. In this way, this chapter addresses our first research objective (Section 1.3): to design a new mechanism based on the CDA that is robust to execution uncertainty, and the second research objective: to study its properties.

4.1 The T-CDA Mechanism

As we pointed out earlier, traditional market mechanisms ignore the execution phase present in every interaction. Given this, here an extension to the CDA[1] is proposed, the T-CDA. Unlike the CDA, the T-CDA allows agents to factor the execution phase into their decision making.

In more detail, the CDA is modified to additionally let agents accept or reject transactions based on the identity of the other agent. To this end, agents not only submit their bids or asks to the market, but also have to explicitly indicate their willingness to interact with a specific agent before a transaction takes place. We call this declaration of willingness a *commitment*. This allows us to leave most of the rules and structure of the CDA intact and also maintains the decentralised nature of the CDA, by leaving the management of trust information and the decision making up to the agents themselves. Indeed, our mechanism does not require agents to reveal this information. As in the CDA, the T-CDA merely provides the necessary means for the agents to communicate their desires effectively. Conversely, this means that agent strategies will be more complex and play an important role in determining individual agent utilities as well as system efficiency, as is the case for the CDA.

In more detail, if $b_i \in B$ has placed a bid o_i^b and $s_j \in S$ has placed an ask o_j^s, we denote the *commitment* of b_i to a transaction based on o_i^b and o_j^s as $c_i(o_i^b, o_j^s)$.

[1] See Section 2.2 for a detailed description of the CDA.

A commitment by s_j would be $c_j(o_i^b, o_j^s)$. Two matching commitments result in a transaction. We do not allow more than one commitment by an agent on its own shout, since there can be only one transaction based on a particular shout. However, we do allow agents to withdraw a commitment, for example because the other agent is not responding. Agents may reject a commitment made by others on their shout. The market maintains a list of all current commitments, in the *commitment book*. The mechanism is defined in detail in Box 4.1, which explains how the T-CDA extends the CDA, as defined in Box 2.1.

We may think of the mechanism as consisting of three components: the bidding and clearing components identified earlier and a new one, the *commitment component*, which manages the interaction with the commitment book, through the commitment accepting rule. This is visualised in Figure 4.1.

To illustrate the trading process, consider a scenario with one buyer, b_0, with $p(b_0) = 1$ and $\ell_0^b = 8$ and one seller, s_1, with $p(s_1) = 0.85$ and $\ell_1^s = 5$. For simplicity, assume both agents have perfect knowledge of $p(\cdot)$. After some bidding, we have the offers $o_0^b = 7$ and $o_1^s = 6.8$ in the order books. Now, in the traditional CDA, the market would immediately clear and a transaction would take place at price $q = 6.9$. However, in the T-CDA, agents consider their expected utility (Equation 3.2) in order to decide whether to commit. It happens that $\hat{u}_1^s(t_{0,1}(6.9)) \geq 0$, so s_1 will commit to $c_1(o_0^b, o_1^s)$. However, $\hat{u}_0^b(t_{0,1}(6.9)) < 0$, so b_0 will reject the commitment, removing it from the commitment book. If s_1 were to improve its ask to $o_1^s = 6.4$, both agents have positive expected utility (at price $q = 6.7$) and they will both commit, resulting in a transaction $t_{0,1}(6.7)$.

4.2 Zero-Intelligence Behaviour

In order to evaluate the mechanism presented above, agent behaviours must also be defined. To this end, this section explains how the ZI strategy can be used in the T-CDA.

The ZI strategy (see Section 2.3.1), although surpassed in terms of efficiency by more modern strategies, remains important for the evaluation of mechanisms, as its uncomplicated behaviour allows the mechanism itself to be investigated, in the absence of the complex effects of a more advanced strategy. Moreover, performance with ZI traders can be said to give a *lower bound* on the expected performance of traders in the mechanism, because ZI agents have practically no intelligence. Therefore, here the ZI strategy for the CDA is adapted to the T-CDA.

Now, in the traditional CDA, an agent's strategy is specified through its bidding behaviour, which dictates the offers an agent submits in the market. In additon to this, a commitment behaviour is also required when trading in the T-CDA, to determine when an agent commits. In order to make informed decisions about when to commit, an agent needs to evaluate the utility it expects to derive from each of the possible transactions. This is given by the estimated expected utility $\tilde{u}$ (Equation 3.4). A rational agent should only commit when the estimated expected utility is non-negative.

Given this, the commitment strategy is based on a single heuristic: if the expected utility is non-negative, an agent a_i is keen on transacting. This heuristic was chosen for its simplicity, easing both analysis and implementation. Other

The T-CDA extends the traditional CDA by separating the bidding and commitment phases implicit in the trading process. Thus, the market no longer clears automatically when a bid and an ask match. Rather, agents themselves must take the initiative in committing to a specific (bid, ask) pair.
In addition to the *order books*, the T-CDA has a *commitment book*, in which a list of all current commitments is maintained. We define an additional rule and adapt the Clearing Rule to deal with commitments:

Commitment Accepting Rule A commitment $c_k(o_i^b, o_j^s)$ is accepted when the prices of the shouts concerned match (i.e. $o_i^b \geq o_j^s$) and one of the shouts was made by the agent committing (i.e. $k = i \vee k = j$). Furthermore, any agent may have only one commitment for a specific shout in the commitment book at any one time. Commitments can be withdrawn by the agent that made them, or rejected by the agent that is being committed to. In either case, the commitment is removed from the commitment book.

Clearing Rule Two commitments match when both the buyer and the seller commit. So, commitments $c_i(o_i^b, o_j^s)$ and $c_j(o_i^b, o_j^s)$ match and would result in a transaction $t_{i,j}(q)$, where q is a transaction price determined by the *Pricing Rule*. After the matching, both the commitments and the shouts concerned are removed from the books.

Box 4.1: The T-CDA protocol is an extension of the CDA protocol (Box 2.1).

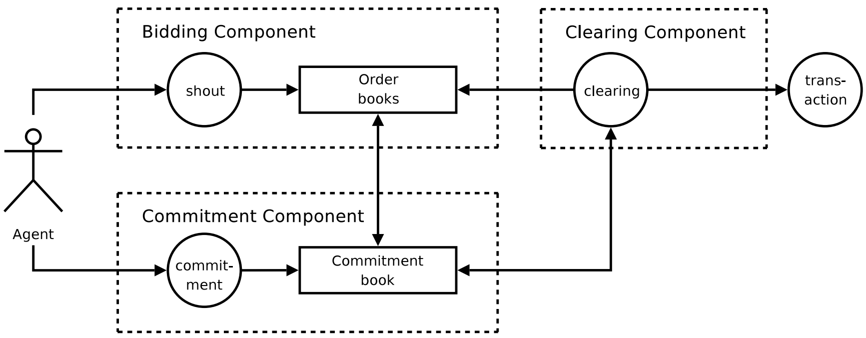

Figure 4.1: Information in the T-CDA flows through three different components. The Commitment component distinguishes the T-CDA from the traditional CDA (Figure 2.3).

strategies may be equally appropriate, e.g. explicitly trading off the utility of transacting now against the possibility of a better transaction later, estimated according to some probability distribution. Hence, the following actions are tried in order:

1. Given commitments to its own shout, a_i picks the best and commits if $\tilde{u} \geq 0$;

2. If a_i is already committed, it does nothing more;

3. Given compatible shouts, a_i picks the best and commits if $\tilde{u} \geq 0$;

4. a_i submits an offer based on the ZI strategy.

If necessary, the agent a_i will withdraw a previous commitment, while any unaccepted commitments on its own shout will be rejected.

In summary, a simple heuristic commitment strategy is added on top of the ZI bidding behaviour, to define a minimally intelligent *lower bound* strategy for the T-CDA. Although this strategy is minimally intelligent, it is rational in the sense that it will not engage in transactions that would result in negative expected utility. In Chapter 5, a more intelligent trading strategy is developed.

4.3 The Simulation

So far, we have defined the T-CDA mechanism as well as a baseline trading strategy for the evaluation of the mechanism. In order to empirically investigate the T-CDA, an environment in which to run experiments is required. An overview of the design and capabilities of the simulator is given in Appendix E. Here, we detail how the trading process is simulated.

In the T-CDA simulator, a *market definition* consists of the definition of several groups of traders (usually two: buyers and sellers). For each group, a strategy, endowment source, execution model and trust source is specified. A *run* of a market definition is subdivided into an arbitrary number of *days*. Every day is subdivided into a pre-specified number of *time steps*. The time step uniquely identifies every moment in simulated time within a run. Hence, if a run consists of n days and m time steps per day, the time step counter ranges from 0 to $nm - 1$.

Every time step, the runner selects a random trader to submit its desired actions (shout, commitments) to the market. The market (in concert with the auctioneer) processes the requests from this trader and attempts to clear the market (i.e. perform as many transactions as possible). The market notifies all traders of any changes. The new information is then processed by the traders and the runner moves the simulation along to the next time step.

Now, there are two exceptions to this normal flow of time in the simulation. First, under specific circumstances, it is possible to determine that no further transactions are possible during the current trading day. In such a case, the runner may stop executing time steps for the current day and move the simulation along to the next day immediately. The time step counter is also incremented to the first step of the next day, to preserve correspondence between time step number and day.

Second, if during a time step an agent a_i commits to a transaction with an agent a_j, the agent a_j is allowed to immediately respond to this by either accepting, rejecting or ignoring the commitment. This is called the Instantaneous Response Step (IRS) assumption, and is optional, but default, behaviour of the simulator. This is done for two reasons:

1. In order to compare the T-CDA to the CDA, it is important that the new mechanism with ZI traders reduces to a normal CDA given

$$p(a_i) = 1 \; ; \; \forall a_i \in A$$

In that way, the impact of one factor (the POS) on the CDA and T-CDA is measured, without confounding the results with differences due to timing.

2. Without the IRS, experiments require a greater number of time steps to run (trials indicate a factor of 10 increase). Therefore, statistically significant results are more readily arrived at when the IRS is enabled.

It is assumed (and verified by trial runs) that the IRS does not have a relevant impact on the eventual efficiency derived by the system. This is supported by the experiments in Section 5.3.3.

4.4 Empirical Evaluation

Now that we have defined the T-CDA, the agent behaviours and the simulation environment, we next detail the empirical evaluation of the T-CDA. In particular we aim to see how it performs with respect to the desiderata specified in Section 3.4. Specifically, we investigate the efficiency of the mechanism and the distribution of utility between buyers and sellers, and the robustness of the mechanism to errors in the trust information. In what follows, we first detail the experimental setup. Then, we detail the results and discuss their implications.

4.4.1 Experiment Settings

This section details the experiment settings. For some variables, although they may impact on the performance of the mechanism in some way, the results obtained here are not sensitive to their specific values. Therefore, for these variables, reasonable default values were chosen (Table 4.1).

More specifically, there are 50 buyer and 50 seller agents. The agents' endowments, which determine the orders the agents have to complete, are generated from a uniform distribution with the range $[6, 8]$ for sellers and $[10, 12]$ for buyers[2]. The maximum price is set to 15. As agents do not learn over trading days (see Section 4.2), a run will consist of a single trading day. Experiments consist of 300 runs per condition. The buyer POS is fixed at 1, because this allows for more insightful analysis, though similar results occur if failure is two-sided.

[2] Although it appears that all traders should transact, this may not be the case, because not all traders may be matched with positive expected utility due to execution uncertainty.

Independent Variables. There are three independent variables. The first two are the expected value $E(pos)$ and variance $Var(pos)$ of the probability of success of sellers. In total, 65 combinations of these variables are run (Table 4.2). If $Var(pos) = 0$, every seller has POS $E(pos)$. Otherwise, POS values are drawn from a Beta distribution with appropriately chosen parameters. The Beta distribution was chosen because it generates values in $[0, 1]$ and allows flexible configuration of μ and σ^2 (see Appendix A for details).

The third variable determines the way in which trust (in sellers) is initialised for the buyers. If trust is CDA-LIKE, a trust of 1 is placed in every seller. This condition thus exhibits the same behaviour as the traditional CDA. With RANDOM trust, trust values are drawn from a uniform distribution. Trust can also be initialised as the MEAN seller POS (i.e. an agent will believe every agent to be as reliable as the population mean reliability), or as a PERFECT copy of the POS value of each seller (i.e. an agent has perfect knowledge of each other agent's POS). Finally, to simulate unreliable trust, the NOISE trust condition initialises trust to the true POS with a certain error value added to it. The trust settings are summarised in Table 4.3.

Metrics. Performance is measured as the sum of the actual (derived) utilities of all buyers, V_B, and the sum of the actual utilities of all sellers, V_S. When the optimal allocation has an expected utility $U \neq 0$, we may express these measures relative to the optimum, as $2V_BU^{-1}$ and $2V_SU^{-1}$, respectively.

Now, we analyse the performance of the mechanism, given that agents have a correct perception of their counterparts' probabilities of success. The analysis serves three main goals. First, it confirms that the emergent behavior of the system is as we expect (Section 4.4.2). Second, we evaluate the behaviour of the mechanism, in comparison to the traditional CDA (Section 4.4.3) and with respect to the optimal performance (Section 4.4.4). Finally, in Section 4.4.5, we evaluate the robustness of the mechanism to errors in the trust information.

4.4.2 Positive Payoff

First of all, calculating the optimal allocation tells us when a positive payoff is possible. We expect that given perfect information, on average, the mechanism will derive a positive utility if that is at all possible.

Hypothesis 1. *If for a certain setting of $E(pos)$ and $Var(pos)$, the optimal buyer utility is positive, so is the expected performance for the* PERFECT *trust setting.*

For the 60 out of 65 combinations of $E(pos)$ and $Var(pos)$ where the optimal expected utility is greater than zero, we do a t-test with the null hypothesis that the mean buyer utility is equal to zero. The alternative hypothesis is that the mean is greater than zero. At the $\alpha = 0.05$ level, we reject the null hypothesis in 56 of the 60 cases.[3]

In the cases where the null hypothesis is not rejected (and the mean buyer utility is thus roughly equal to zero), the estimated mean is greater than zero,

[3] If we protect the null hypothesis against spurious results by setting $\alpha' = 1 - 0.95^{1/65}$ (Cohen, 1995), the null hypothesis is rejected in 54 cases.

Variable	Value	Unit
# of buyers	50	agents
# of sellers	50	agents
Maximum price	15	-
Seller endowments	$[6, 8]$	-
Buyer endowments	$[10, 12]$	-

Table 4.1: Fixed variable values

	$Var(pos)$									
$E(pos)$	0.000	0.010	0.045	0.085	0.125	0.155	0.185	0.205	0.235	0.245
0.10	✓	✓	✓	✓						
0.30	✓	✓	✓	✓	✓	✓	✓	✓		
0.50	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
0.60	✓	✓	✓	✓	✓	✓	✓	✓	✓	
0.70	✓	✓	✓	✓	✓	✓	✓	✓		
0.75	✓	✓	✓	✓	✓	✓	✓			
0.80	✓	✓	✓	✓	✓	✓				
0.85	✓	✓	✓	✓	✓					
0.90	✓	✓	✓	✓						
0.95	✓	✓	✓							
1.00	✓									

Table 4.2: Levels of $E(pos)$ and $Var(pos)$. A checkmark in the table indicates a combination of $E(pos)$ and $Var(pos)$ that has been run.

condition	initialisation
NAIVE	$trust_i(a_j) = 1$
RANDOM	$trust_i(a_j) = U(0, 1)$
MEAN	$trust_i(a_j) = \frac{1}{\lvert S \rvert} \sum_{s_k \in S} p_s(s_k)$
PERFECT	$trust_i(a_j) = p_s(a_j)$
NOISE	$trust_i(a_j) = p_s(a_j) + N(0, x)$, bounded to $[0, 1]$
	$x \in \{0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50\}$

Table 4.3: Trust levels. U is the Uniform distribution. N is the Normal (Gaussian) distribution. The NOISE condition is introduced in Section 4.4.5.

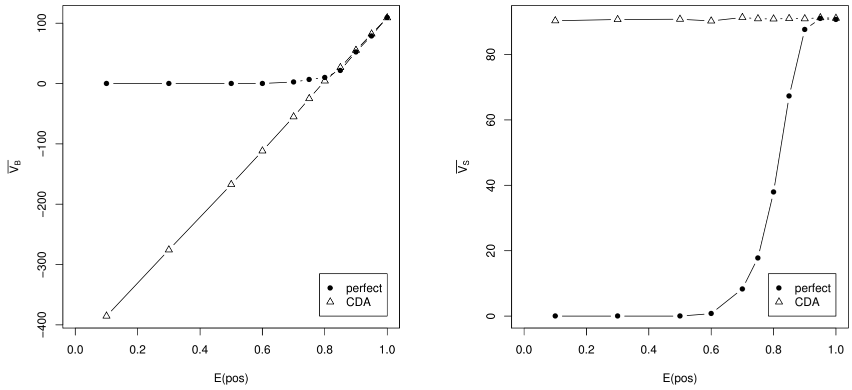

(a) Buyer utility (b) Seller utility

Figure 4.2: Buyer and seller utility for the Trust-based CDA given PERFECT trust and the normal CDA. In (a), PERFECT trust avoids making a loss, where the CDA does make a loss. Panel (b) shows that the normal CDA allows unreliable sellers to exploit buyers.

so we need not consider the alternative that the actual mean is smaller than zero. Furthermore, these cases all have a very small optimal expected utility. Hence, in general, the mechanism does derive a positive expected utility if this is possible.

4.4.3 Comparison to the CDA

It has been shown that, given PERFECT trust information, ZI traders in the T-CDA avoid making a loss and turn a profit whenever possible. Now, we also show that they do better than agents that do not take trust information into account, as in the CDA-LIKE and the RANDOM conditions.

To this end, Figure 4.2(a) shows a typical outcome when $Var(pos) = 0$, for different levels of $E(pos)$. The PERFECT condition does not trade for low values of $E(pos)$, where a profit is not possible. For higher values of $E(pos)$, the utility for the PERFECT condition increases more or less linearly. For the CDA-LIKE condition, the relationship between $E(pos)$ and buyer utility is linear, which is what we expect, since it will ignore the probability of success of sellers altogether. Hence, it derives a (very large) negative utility for low $E(pos)$. Beyond a certain threshold, there is very little distinction between the PERFECT and CDA-LIKE conditions. This is to be expected, since then transactions are usually desirable and given random (ZI) bidding both conditions will lead to approximately identical results.

In Figure 4.2(b) we see that the influence of $E(pos)$ on seller utility is quite different. Clearly, accurate trust information prevents the buyers from being exploited by sellers. We return to this point later, in Section 4.4.4.

The above conceptions are formalised as follows:

Hypothesis 2. *Under any setting of $E(pos)$ and $Var(pos)$,* PERFECT *trust will* *do at least as well (in terms of buyer utility) as the* RANDOM*,* CDA-LIKE *and* MEAN *conditions.*

To test this hypothesis, for all combinations of $E(pos)$ and $Var(pos)$, pairwise comparisons of the PERFECT condition were done against the other conditions. Two t-tests were performed for each pair, in both cases the null hypothesis is that the means are equal. In the first test, the alternative is that the mean in the PERFECT condition is greater, in the second, that the mean is less. The resulting p-values were inspected at $\alpha' = 1 - 0.95^{1/65}$, protecting the null hypothesis (no difference) against spurious results. Here, the important results are given. A comprehensive view of the data is provided in Appendix B.

Note that given the experiment settings, when $Var(pos) = 0$ and $E(pos) \geq 0.82$, the decisions made by CDA-LIKE trust are, on average, rational.[4] Hence, we cannot expect much advantage from good trust information in that case. In comparison to MEAN, we expect no difference when $Var(pos) = 0$. Also see Figure 4.4, Table 4.4 and the corresponding discussion, which show that for $E(pos) \geq 0.80$ and $Var(pos) = 0$, errors in the trust information have very little impact on the overall system performance.

Looking at the 'PERFECT > other' alternative hypothesis, at α', PERFECT is significantly better than RANDOM in 64 of the 65 cases, better than CDA-LIKE in 57 of the 65 cases and better than MEAN in 43 of the 65 cases. The cases of no difference correspond to the expectations mentioned above. For the 'PERFECT < other' alternative, there are no significant differences at α'.

Thus, it is safe to say that the PERFECT condition improves upon the control conditions RANDOM, CDA-LIKE and MEAN. Moreover, it is clear that the T-CDA does better than the CDA when faced with uncertainty about the result of transactions.

4.4.4 Benchmark

In this experiment we benchmark the T-CDA's performance against the optimal performance and make some overall qualitative observations about its behaviour.

To this end, Figure 4.3(a) shows the total utility achieved by the system, normalised by the maximum expected utility from the optimal allocation. The mechanism does well when either $Var(pos)$ is high, or $E(pos)$ is high, or both. This is because, in both cases, the part of the population from which profit can be derived have $E(pos) \approx 1$. Hence, when buyers bid randomly from $[0, \ell]$, they are submitting profitable bids. If, however, a large group from which profit may potentially be derived has a low POS, the bidding strategy does poorly. This is because it submits bids that are too high (overbidding). Hence, the agent itself is not willing to transact at that price, given the execution uncertainty. This means that the number of transactions that occur is reduced, in turn reducing total utility. Hence, the figure reveals the need for an agent's bidding strategy to be informed by its trust function in order to bid appropriately.

Another relevant aspect of the behaviour is the balance of utility between buyers and sellers. This is shown in Figure 4.3(b). In the $Var(pos) = 0$ case, it

[4] Assume the transaction price is, on average, the equilibrium price $\bar{q} = 9$. Then, given the average limit price for buyers, $\bar{\ell^b} = 11$ and that all sellers have the same POS p, we can find p such that expected buyer utility (Equation 3.2), on average, is non-negative: $u^b = \bar{\ell^b} p - \bar{q} \geq 0 \Rightarrow p \geq \frac{9}{11} \approx 0.82$.

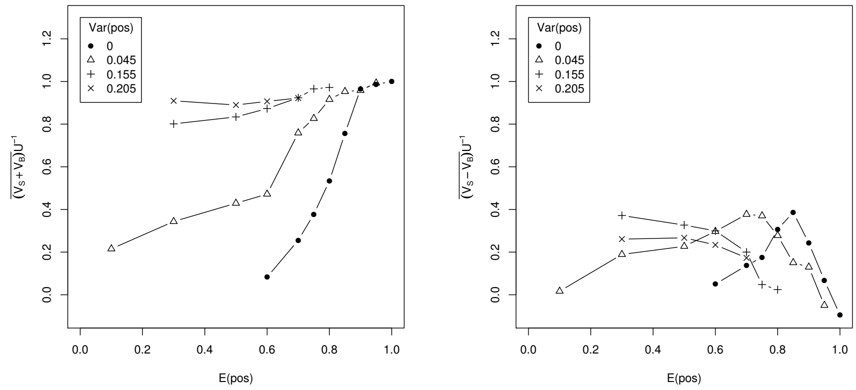

(a) Normalised total utility

(b) Normalised difference

Figure 4.3: The normalised utility or *efficiency* derived by the mechanism and the disparity between seller and buyer utilities.

appears that sellers are the first to profit from an increase in $E(pos)$, with the balance being restored only for the highest values of $E(pos)$. Specifically, for $E(pos) = 0.60$, observe that the difference of seller and buyer utility is almost identical to the total utility in the system, i.e. only the sellers turn a significant profit. The higher $Var(pos)$ levels show an imbalance that decreases when $E(pos)$ increases. Once again, the imbalance is caused by the bidding strategy, which is uninformed about the actual worth of the sellers' offers.

4.4.5 Effect of Noise

Now, we analyse the effect of the degradation of trust information on the mechanism. To simulate unreliable trust information, each buyer's trust function is initialised to the actual POS values with some arbitrary level of Gaussian noise applied to it. Figure 4.4 provides an overview of the results.

The figure provides a number of interesting insights. First, if the noise level is high, performance degrades almost linearly with $E(pos)$. This is to be expected, since interaction partners are chosen almost completely at random, and this randomness leads to a linear relationship between buyer utility and $E(pos)$. Second, if $E(pos)$ is very low, performance increases linearly with a decreasing noise level, until a 'plateau' is reached where utility is zero. A linear regression (Table 4.4) shows that a linear relation can indeed account for a large proportion of the variance in these cases. Adding noise means that agents will overestimate POS in some cases and hence that they may transact even if it is not in their best interest, leading to losses. The 'plateau' where utility is zero exists because even with some overestimation of the POS, agents do not see transactions as desirable.

Last, when $E(pos) \geq 0.80$, the noise level seems to have very little impact on the total utility derived by buyers, rather increasing linearly with an increasing $E(pos)$. Linear regression of buyer utility on noise (Table 4.4) confirms this. This may be explained by the fact that in Figure 4.4, $Var(pos) = 0$ and hence

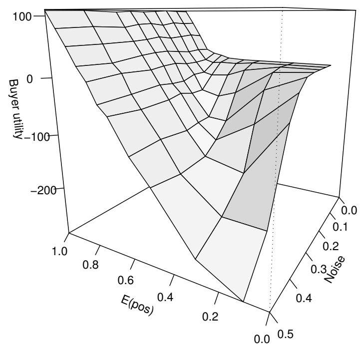

Figure 4.4: Performance degrades when $E(pos)$ is lowered and when more noise is added to trust values. The noise level represents the variance of the Gaussian noise distribution that is applied to individual agents' trust function.

$E(pos)$	$noise$	r^2	F	p
0.10	$[0.15, 0.50]$	0.74	280	$\ll 0.01$
0.30	$[0.10, 0.50]$	0.70	226	$\ll 0.01$
0.50	$[0.05, 0.25]$	0.61	155	$\ll 0.01$
0.80	$[0.00, 0.50]$	0.00	0.15	> 0.50
0.85	$[0.00, 0.50]$	0.00	0.04	> 0.50
0.90	$[0.00, 0.50]$	0.00	0.01	> 0.50
0.95	$[0.00, 0.50]$	0.00	0.01	> 0.50
1.00	$[0.00, 0.50]$	0.00	0.00	> 0.50

Table 4.4: Linear regression of buyer utility on noise, for $Var(pos) = 0$, significance tested against F distribution. r^2 is the proportion of the total variance accounted for by the regression line. F is the value of the F-test statistic for a linear regression and p is the significance of the regression line given by an F distribution.

there is no benefit in distinguishing between sellers. The intuition behind this is that the application of noise introduces an arbitrary preference for certain sellers, which is different for each buyer, and transactions are usually desirable. Thus, the effects of noise on the individual cancel out over the entire population.

4.5 Summary

This chapter has provided an overview of the T-CDA mechanism, which allows traders to take execution uncertainty into account in their decision making. Specifically, it allows agents to use the trust they place in each other in order to decide which potential transaction to commit to. The decentralised nature of the CDA is preserved, by requiring traders to take the initiative in committing to a transaction based on a bid and an ask previously placed in the market. Moreover, the ZI bidding behaviour for the CDA is adapted to the T-CDA by adding a simple heuristic commitment behaviour on top of the unmodified bidding behaviour.

Subsequently, the T-CDA simulator that is used for the empirical investigation of the new mechanism has been described and, specifically, the way time is simulated has been discussed. Then, the mechanism has been empirically investigated with ZI traders. It has been demonstrated to be robust against increasing execution uncertainty. The CDA on the other hand, is shown to break down. Moreover, the effect of unreliability of trust information on market efficiency is shown to be linear.

Thus, this chapter has addressed the first two research objectives (Section 1.3), to develop an extension of the CDA that is robust to execution uncertainty and to study its properties. However, although unlike the CDA, the T-CDA does not derive negative social welfare when execution uncertainty is high, market efficiency in the T-CDA decreases with increasing execution uncertainty. It was hypothesised that this is because the bidding range of ZI agents is inappropriate. To show that this is indeed the case, in the next chapter, a rudimentary trading strategy based on ZIP is developed. Moreover, its evaluation provides further insights into the T-CDA mechanism and the challenges in designing trading strategies for this new mechanism.

Chapter 5

Designing a Trading Strategy

This chapter details the Trust-Based ZIP (T-ZIP) strategy. T-ZIP is a rudimentary trading strategy for the T-CDA. The main aim in developing this strategy is to show that the inefficiency of the T-CDA with ZI agents when the probability of success is low, is indeed a result of the inappropriate bidding range of the ZI agents rather than an inherent weakness of the T-CDA mechanism. Therefore, the T-ZIP is a straightforward adaptation of the ZIP strategy (see Section 2.3.2) and it is not designed with all possibilities in mind. Specifically, it is assumed that although in general there is no single equilibrium price (Section 3.2), there is a single equivalent price (see Section 5.1.1) that traders should converge to.

Thus, the T-ZIP strategy serves as a further demonstration of the capabilities of the T-CDA mechanism and as a starting point for good trading strategies, but not as a generally applicable trading strategy in its own right. Within the context of the research objectives of Section 1.3, this Chapter mainly addresses objectives three and four, concerning the development and analysis of a trading strategy of the T-CDA, but also furthers the fulfillment of the second objective, the study of the properties of the T-CDA mechanism.

The T-ZIP strategy is detailed in Section 5.1. Its implementation is discussed in Section 5.2 and it is empirically evaluated in Section 5.3. A discussion of the results and recommendations for trading strategies are given in Section 5.4. Finally, Section 5.5 summarises the main points.

5.1 The Trust-Based ZIP Strategy

A trader a_i employing the ZIP strategy uses bid and ask prices to determine a *target price* $\tau_i(t)$. Then, it applies a machine learning algorithm (adaptation mechanism) that updates the *profit margin* μ_i of the trader by adjusting it so that its shout price q_i (Equation 2.2) goes towards $\tau_i(t)$. For a detailed description of ZIP, refer back to Section 2.3.2.

The ZIP strategy cannot be applied directly to the T-CDA, for the following reasons:

1. In the T-CDA, shout prices cannot be taken at face value, because execution uncertainty needs to be taken into account when determining the value of a shout;
2. Because of the introduction of the commitment step, learning opportunities in the T-CDA differ from those in the CDA;
3. As noted previously, ZIP works under the assumption that there is an equilibrium price it should converge to. This is not always the case in the T-CDA.

Now, when developing the T-ZIP strategy (as mentioned above), we choose to assume that there is a price an agent should converge to. Thus, in designing T-ZIP, the first two problems must be addressed. First, we need to define how the target price $\tau_i(t)$ is set and, then, we need to re-define the rules that determine when and how to update the profit margin. The adaptation mechanism does not need to be modified.

5.1.1 Setting the Target Price

In the ZIP strategy, the profit margin is adjusted to approach a target price $\tau_i(t)$. The target price $\tau_i(t)$ is calculated from a shout price or transaction price $q(t)$ by applying small random perturbations (Equation 2.7). However, in the T-CDA, we need to take Execution Uncertainty into account. Specifically, if we wish to update the profit margin based on a transaction price $q(t)$, we need to take the POS of the traders involved into account. For example, if a buyer b_i with POS 1 wants to adjust its margin based on a transaction price $q(t)$, and if the buyer b_j involved in the transaction also has POS 1, the target price can simply be based on $q(t)$. However, if b_j has POS 0.8, then b_i could get away with bidding a lower price, since we assume sellers make their choices based on the expected utility of an offer.

Therefore, an agent a_i should calculate an adjusted price $\upsilon_i(t)$ that represents the *equivalent value* of a shout or transaction price relative to its own POS. The target price $\tau_i(t)$ may then be redefined based on the equivalent value $\upsilon_i(t)$:

$$\tau_i(t) = \mathcal{R}_i(t)\upsilon_i(t) + \mathcal{A}_i(t) \ , \tag{5.1}$$

where $\mathcal{R}_i$ and $\mathcal{A}_i$ are the random perturbations defined in Section 2.3.2. Compare this equation to Equation 2.7.

In the following, the equivalent value $\upsilon_i(t)$ is derived from the expected utility function $\hat{u}_i$ (Equation 3.2). Buyers and sellers are considered separately. Note that in defining $\upsilon_i(t)$, the POS function $p(\cdot)$ is used. In practical implementations, agents will substitute their trust function for the actual POS function.

Buyers Now, we consider how a buyer b_i should calculate the equivalent value of a certain price $q(t)$. Say that $q(t)$ was shouted by a buyer b_j, or that b_j transacted at price $q(t)$. Now, we must determine what value b_i would need to shout to make an offer of the same value (expected utility) to a seller s_k. This can be expressed as the following equation:

$$\hat{u}^s_k(t_{i,k}(\upsilon_i(t))) = \hat{u}^s_k(t_{j,k}(q(t))) \ . \tag{5.2}$$

Here, $\upsilon_i(t)$ is the price at which b_i would transact with s_k, such that s_k derives the same expected utility from this transaction as it would derive from a transaction with b_j at price $q(t)$. When substituting the definition of expected utility (Equation 3.2) into this equation, we obtain the following equality:

$$\upsilon_i(t)p(b_i) - \ell_k^s = q(t)p(b_j) - \ell_k^s \ , \tag{5.3}$$

which reduces to the following solution for $\upsilon_i(t)$:

$$\upsilon_i(t) = q(t)\frac{p(b_j)}{p(b_i)} \ . \tag{5.4}$$

Hence, to find the equivalent price of $q(t)$, we normalise it with the ratio of the competing buyer's POS to the agent's own POS. Intuitively, this is because sellers scale the transaction price according to (their estimate of) the buyer's POS when determining the utility of a transaction. Thus, if the other buyer has a higher POS, we would need to bid a higher value, which is what one would expect. Note that the solution is independent of the identity of the seller s_k.

Sellers We have defined equivalent prices for buyers. Now, we consider how a seller s_i should calculate the equivalent value of a price $q(t)$. Say $q(t)$ was shouted by a seller s_j, or s_j transacted at price $q(t)$. We determine what value seller s_i would need to shout to make an offer of the same value to a buyer b_k. Again, using expected utility (Equation 3.2), we derive the following equation:

$$\hat{u}_k^b(t_{k,i}(\upsilon_i(t))) = \hat{u}_k^b(t_{k,j}(q(t))) \tag{5.5}$$

$$\ell_k^b p(s_i) - \upsilon_i(t) = \ell_k^b p(s_j) - q(t) \ . \tag{5.6}$$

Again, $\upsilon_i(t)$ is the price at which seller s_i would transact with b_k, such that b_k derives equal expected utility from this transaction as from a transaction with b_j at price $q(t)$. Solving for $\upsilon_i(t)$:

$$-\upsilon_i(t) = \ell_k^b p(s_j) - q(t) - \ell_k^b p(s_i) \tag{5.7}$$

$$\upsilon_i(t) = \ell_k^b(p(s_i) - p(s_j)) + q(t) \ . \tag{5.8}$$

This solution is strikingly different from Equation 5.4, because here $\upsilon_i(t)$ is proportional to the difference of the agent's POS with the POS of the competing seller, not the ratio between the two. Moreover, this solution is problematic, as the equivalent value depends on the limit price of a buyer. This, of course, may vary from buyer to buyer and hence there is no single equivalent value. The intuition behind this is that, unlike sellers, buyers do not scale the transaction price according to (their estimate of) the POS of the seller, but adjust their limit price accordingly. Hence, the relationship between the transaction price and the equivalent price is not as straightforward as for sellers.

We do know that $\ell_k^b > 0$ and thus, the sign of the difference $p(s_i) - p(s_j)$ will tell us whether $\upsilon_i(t) > q(t)$, $\upsilon_i(t) = q(t)$ or $\upsilon_i(t) < q(t)$. However, we do not know the relationship between $\upsilon_i(t)$ and q_i, the price s_i would currently shout in the absence of this information.

Now, we cannot make any assumptions on ℓ_k^b, except that it has a certain relation to $q(t)$. Therefore, to simplify Equation 5.8, we introduce a parameter κ_k that expresses this relationship:

$$\ell_k^b = q(t) \cdot \kappa_k(t) \ . \tag{5.9}$$

Now, κ_k may be thought of as the *inverse* profit margin, by rewriting the previous equation as follows:

$$\frac{1}{\kappa_k(t)} = \frac{q(t)}{\ell_k^b} = 1 + \mu_k \ , \tag{5.10}$$

where the last equality is due to Equation 2.2. From Equation 5.10 and the limits on μ_k in Section 2.3.2, we know that for buyers, $\mu_k \in [-1, 0]$, so $\frac{1}{\kappa_k(t)} \in [0, 1]$ and hence $\kappa_k \in [1, \infty)$.

The parameter κ_k allows us to eliminate the limit price ℓ_k^b from the equivalent price formula, by substituting Equation 5.9 into Equation 5.8:

$$\upsilon_i(t) = (q(t)\kappa_k(t))(p(s_i) - p(s_j)) + q(t) \tag{5.11}$$
$$= q(t)(\kappa_k(t)(p(s_i) - p(s_j)) + 1) \ . \tag{5.12}$$

It seems that we have gained little by this transformation: this formula still contains an unknown. However, κ_k abstracts away the price level in the specific market. Thus, we can more easily come up with reasonable assumptions on κ_k, without knowledge of the specific market.

Now, from Algorithm 2.1, we know that based on a competing ask, we can be asked to either raise or lower the profit margin. For the sake of brevity, define $d = p(s_i) - p(s_j)$ and note that $d \in [-1, 1]$. In order to decide whether to raise or lower the margin, we need to decide the following inequalities:

$$q_i \leq \upsilon_i(q) = q(t)(\kappa_k d + 1) \tag{5.13}$$
$$q_i \geq \upsilon_i(q) = q(t)(\kappa_k d + 1) \tag{5.14}$$

These are the analogues of $q_i \leq q$ and $q_i \geq q$ in Algorithm 2.1, respectively. Now, because we do not know κ_k, we cannot do this exactly. Therefore, an approximate answer is required. Because we do not want to raise or lower the margin unnecessarily, we define two conservative estimates: $\upsilon_i^\uparrow$ for when the margin should be raised and $\upsilon_i^\downarrow$ for when the margin should be lowered. To be conservative (i.e. prevent unwarranted adaptation of the margin), these estimates should satisfy the following constraints:

$$\textbf{if } q_i \leq \upsilon_i^\uparrow(q) \textbf{ then } q_i \leq \upsilon_i(q) \tag{5.15}$$
$$\textbf{if } q_i \geq \upsilon_i^\downarrow(q) \textbf{ then } q_i \geq \upsilon_i(q) \tag{5.16}$$

Assume we have bounds $\kappa_{\min}$ and $\kappa_{\max}$ for κ_k that satisfy:

$$\kappa_{\min} \leq \kappa_k \leq \kappa_{\max} \ ; \ \forall k \tag{5.17}$$

Then we can define the estimates $\upsilon_i^\uparrow$ and $\upsilon_i^\downarrow$ as follows:

$$\upsilon_i^\uparrow(q) = \begin{cases} q(t)(\kappa_{\min} d + 1) & d \geq 0 \\ q(t)(\kappa_{\max} d + 1) & d < 0 \end{cases} \tag{5.18}$$

$$\upsilon_i^\downarrow(q) = \begin{cases} q(t)(\kappa_{\max} d + 1) & d \geq 0 \\ q(t)(\kappa_{\min} d + 1) & d < 0 \end{cases} \tag{5.19}$$

Now, it must be shown that these functions are indeed conservative estimates of $\upsilon_i(q)$, as defined by Equation 5.15 and Equation 5.16.

Theorem 5.1. *The definition in Equation 5.18 satisfies the constraint given in Equation 5.15, given that Equation 5.17 holds. Specifically, for all* $d \in [-1, 1]$ *and* $q(t) > 0$, **if** $q_i \leq \upsilon_i^{\uparrow}(q)$ **then** $q_i \leq \upsilon_i(q)$.

Proof. Note that to show that Equation 5.15 holds, it suffices to show that:

$$\upsilon_i^{\uparrow}(q) \leq \upsilon_i(q) \tag{5.20}$$

Case I: $d \geq 0$. In this case, Equation 5.18 defines $\upsilon_i^{\uparrow}$ as:

$$\upsilon_i^{\downarrow}(q) = q(t)(\kappa_{\min} d + 1) \tag{5.21}$$

By substituting Equation 5.21 and Equation 5.11 into Equation 5.20, we obtain:

$$\upsilon_i^{\uparrow}(q) \leq \upsilon_i(q) \tag{5.22}$$
$$q(t)(\kappa_{\min} d + 1) \leq q(t)(\kappa_k d + 1) \tag{5.23}$$
$$\kappa_{\min} d + 1 \leq \kappa_k d + 1 \tag{5.24}$$
$$\kappa_{\min} d \leq \kappa_k d \tag{5.25}$$

Now, because $d \geq 0$, this reduces to:

$$\kappa_{\min} \leq \kappa_k \tag{5.26}$$

This holds because of Equation 5.17. Thus, when $d \geq 0$, the theorem holds. Now, it remains to show the same for $d < 0$.

Case II: $d < 0$. In this case, Equation 5.18 defines $\upsilon_i^{\uparrow}$ as:

$$\upsilon_i^{\uparrow}(q) = q(t)(\kappa_{\max} d + 1) \tag{5.27}$$

Again, substituting the relevant definitions into Equation 5.20 gives:

$$\upsilon_i^{\uparrow}(q) \leq \upsilon_i(q) \tag{5.28}$$
$$q(t)(\kappa_{\max} d + 1) \leq q(t)(\kappa_k d + 1) \tag{5.29}$$
$$\kappa_{\max} d \leq \kappa_k d \tag{5.30}$$

And because $d < 0$, this reduces to:

$$\kappa_{\max} \geq \kappa_k \tag{5.31}$$

This holds according to Equation 5.17. □

Theorem 5.2. *The definition in Equation 5.19 satisfies the constraint given in Equation 5.16, given that Equation 5.17 holds. Specifically, for all* $d \in [-1, 1]$ *and* $q(t) > 0$, **if** $q_i \geq \upsilon_i^{\downarrow}(q)$ **then** $q_i \geq \upsilon_i(q)$.

Proof. Analogous to the proof of the previous theorem. □

Therefore, Equation 5.18 and Equation 5.19 indeed provide conservative estimates of $\upsilon_i(q)$. These estimates can be used both to decide when to raise or lower the margin and to set the target price $\tau_i(t)$. Note that for buyers, we can define these functions as $\upsilon_i^{\uparrow}(t) = \upsilon_i^{\downarrow}(t) = \upsilon_i(t)$.

For now, let $\kappa_{\min}$ and $\kappa_{\max}$ be parameters to the algorithm. Note that they have no impact on υ_i in the case that $d = 0$, i.e. when $Var(pos : S) = 0$. A more advanced strategy could attempt to estimate these values with more accuracy and on an agent-by-agent basis.

5.1.2 Bargaining Strategy

In the ZIP bargaining mechanism, an offer successfully being accepted as a transaction or not is the input for the learning mechanism. If a transaction occurs, this is interpreted as positive feedback. If it does not, this is considered negative feedback. However, given that we additionally have commitments, these events cannot be interpreted in the same way. Specifically, the fact that an offer is submitted and no transaction immediately results does not necessarily constitute a negative response to that bid. On the other hand, when a commitment is made to a bid, this is unmistakably a positive reaction. Similarly, a rejection of a commitment is a negative response. Therefore, we consider commitments as a learning opportunity.

In addition to this, we might encounter a situation in which no commitments are being made, perhaps because at the current prices, no commitments are desirable. Hence, we also need to consider lowering the margin when no commitments are being made, to improve the chances of transacting.

Commitment-based Bargaining

Now, we consider how an agent can update its profit margin when a commitment is made in the market. Before discussing the learning rules in detail, we introduce a general framework that clearly shows the extension of the ZIP rules to the T-CDA.

Specifically, we generalise the bargaining mechanism given in Algorithm 2.2 and Algorithm 2.1, such that it is valid for both buyers and sellers (Algorithm 5.1). First, ZIP defines two learning opportunities: 'the last shout resulted in a transaction at price q' is considered as positive feedback on the price q and 'the last shout, at price q, did not result in a transaction' is considered as negative feedback. These rules are denoted by $\mathbf{event}_\oplus$ and $\mathbf{event}_\ominus$, respectively.

Furthermore, there are rules that determine when to raise or lower the margins based on this feedback. In response to positive feedback, the margin can be either raised or lowered. After a negative feedback, the margin can only be lowered. The rules that determine whether to do this are $\mathbf{condR}_\oplus$, $\mathbf{condL}_\oplus$ and $\mathbf{condL}_\ominus$. Finally, the previous section has shown that when lowering and raising the margin, the appropriate target price differs. Therefore, $\upsilon_i^\downarrow(q)$ and $\upsilon_i^\uparrow(q)$ are specified explicitly as targets in Algorithm 5.1. For ZIP, $\upsilon_i^\downarrow(q) = \upsilon_i^\uparrow(q) = q$.

Now, we discuss these five rules in turn. This consists of the rule as adopted by ZIP and the rationale behind it. Then, the new rule for T-ZIP is derived.

$\mathbf{event}_\oplus$ In the ZIP strategy, a transaction occurring immediately after a shout is submitted is considered positive feedback. As was noted before, in this section we consider a commitment on a shout as positive feedback. In more detail, a commitment by an agent a_i to transact with agent a_j at price q is considered positive feedback on the price q.

$\mathbf{event}_\ominus$ Originally, no transaction resulting from a shout is considered negative feedback. However, in the T-CDA, because agents must take the initiative in committing to transactions, transactions never immediately result from shouts. Therefore, we cannot always consider this as negative feedback. Here,

Algorithm 5.1 Revised bargaining algorithm

if event$_\oplus$ **then**
  if condR$_\oplus$ is met **then**
    raise profit margin towards $\upsilon_i^\uparrow(q)$
  end if
  if condL$_\oplus$ is met **then**
    lower profit margin towards $\upsilon_i^\downarrow(q)$
  end if
else if event$_\ominus$ **then**
  if condL$_\ominus$ is met **then**
    lower profit margin towards $\upsilon_i^\downarrow(q)$
  end if
end if

we consider negative feedback based on commitments (negative feedback based on shouts will be considered later). If an agent a_j has committed to a transaction with a_i at price q and if a_i rejects this commitment by a_j, that is considered negative feedback on the price q.

condR$_\oplus$ A trader can raise its margin regardless of whether it is active or not. The intuition behind raising the margin after positive feedback is that positive feedback to a certain price q must mean that q is a competitive price. Hence, if bidding q would mean a higher margin to an agent a_i, that means a_i can raise its margin. For example, a ZIP buyer would raise its margin if $q_i \geq \upsilon_i^\uparrow(q)$, i.e. if it would have bid a higher value than $\upsilon_i^\uparrow(q)$. However, a commitment in the T-CDA is one-sided feedback (i.e. a commitment by a buyer to transact with a specific seller does not give us any information about the seller's willingness to transact with that buyer), unlike a transaction in the CDA, so T-ZIP buyers should only learn based on feedback by sellers and vice versa.

a. **Buyers**: last commit (price q) is by a seller and $q_i \geq \upsilon_i^\uparrow(q)$

b. **Sellers**: last commit (price q) is by a buyer and $q_i \leq \upsilon_i^\uparrow(q)$

condL$_\oplus$ The margin should only be lowered when an agent is active, because if it is inactive, this means it has already successfully filled all its orders. Hence, the agent has no incentive to lower its margin. For a seller, if there is a positive reaction to an ask that would imply a lower margin, this means that the seller risks being undercut by the competition. Thus, the seller should lower its margin. The same holds for buyers and competing bids.

The risk of being undercut by the competition exists for agents in the T-CDA as well. However, they must deal with an additional problem, namely, that they can lower their margin too much. That is, the limit price alone does not fully determine the price for which $\hat{u}_i \geq 0$, rather it is also determined by the POS of the trading partner. Hence, the margin might be lowered to such a point that the agent itself is not willing to trade at that price. E.g. a seller s_j is at risk of being undercut by another seller s_k, because b_i has committed to a transaction $t_{i,k}(q)$. However, if s_j were to transact at the equivalent price, this would result in negative utility: $\tilde{u}_j^s(t_{i,j}(\upsilon_i^{\downarrow}(q))) < 0$. Hence, s_j should not lower its margin.[1] Thus, we arrive at the following rules:

a. **Buyers**: last commit (price q) is by a buyer AND $q_i \leq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^b(t_{i,j}(\upsilon_i^{\downarrow}(q))) \geq 0$; s_j the seller involved in last commit

b. **Sellers**: last commit (price q) is by a seller AND $q_i \geq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^s(t_{j,i}(\upsilon_i^{\downarrow}(q))) \geq 0$; b_j the buyer involved in last commit

condL$_\ominus$ As noted previously (see **condL**$_\oplus$), an agent should only lower its margin when it is still active in the market. In the original ZIP, a seller would lower its margin when there is a negative response to an offer with price $q < q_i$, since an offer it would make at q_i would similarly be rejected. The same basic reasoning holds for the T-ZIP as well, though we now compare q_i to the equivalent price $\upsilon_i^{\downarrow}(q)$. The risk of lowering the margins too much, identified for **condL**$_\oplus$, holds here as well. Hence, the new rules are as follows:

a. **Buyers**: last rejection (price q) is by a seller AND $q_i \leq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^b(t_{i,j}(\upsilon_i^{\downarrow}(q))) \geq 0$; s_j the rejected seller

b. **Sellers**: last rejection (price q) is by a buyer AND $q_i \geq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^s(t_{j,i}(\upsilon_i^{\downarrow}(q))) \geq 0$; b_j the rejected buyer

Shout-based Bargaining

Having designed the rules for commitment-based learning, we now turn to shout-based bargaining. Shout-based bargaining is required because in certain market conditions, it may be the case that bids and asks are being submitted, but no commitments are yet being made, because the bid-ask spread does not allow it, or traders do not expect positive utility from a transaction at current prices. Therefore, it is necessary to lower the margin based on shouts being submitted to the market to ensure that trade actually occurs. Note that in the T-CDA, shout-based bargaining is always negative feedback. Positive feedback is always a commitment being made.

The question, then, is: *when* is it appropriate to lower a trader's margin based on competing shouts being submitted? Here, the approach we adopt is to view shouts that are being submitted as an information source inferior to commitment information. Hence, for a seller, when there currently are commitments on asks by other sellers, there is no need to lower the margin based on new asks being submitted. Given this additional source of negative feedback, we extend the rule **event**$_\ominus$ with the following:

a. **Buyers**: A shout was submitted, and the commitment book contains no commitments by sellers.

b. **Sellers**: A shout was submitted, and the commitment book contains no commitments by buyers.

[1] Agent s_j should not lower its margin to $\upsilon_i^{\downarrow}(q)$. However, a scheme could be devised that sets an alternative target price. Several approaches were tried and the one adopted here seems to work well (i.e. not lowering the margin at all).

Now, we must also define how to decide whether to update the margin. First of all, the margin should only be updated according to *competing* shouts (for buyers, based on bids; for sellers, based on asks), because our motivation for lowering the margin is to avoid being undercut by the competition. Furthermore, we must determine whether a transaction at a price equivalent to our current price would be desirable. However, unlike in the case of a rejection, it is not clear which agent the transaction would be with. Thus, we must pick an agent to evaluate this against. The approach chosen here[2] is to evaluate against the agent with the highest POS of all agents of the opposing type that currently have a shout in the order book.

Based on this, a rule can be defined, that determines when to lower the margin based on the shout-based negative feedback. $\mathbf{condL}_{\ominus}$ is extended by the following:

a. **Buyers**: last shout was a bid AND $q_i \leq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^b(t_{i,j}(\upsilon_i^{\downarrow}(q))) \geq 0$; s_j best seller in order book

b. **Sellers**: last shout was an ask AND $q_i \geq \upsilon_i^{\downarrow}(q)$ AND $\tilde{u}_i^s(t_{j,i}(\upsilon_i^{\downarrow}(q))) \geq 0$; b_j best buyer in order book

This concludes the description of the T-ZIP trading strategy for the T-CDA. A method was developed to set the target price for the adaptation mechanism, taking the different estimated POS of different agents into account. Next, the rationale behind the ZIP bargaining rules was applied to the T-CDA in order to derive the bargaining rules for T-ZIP. It must be noted that T-ZIP does not yet constitute a generally applicable trading strategy for the T-CDA, as the assumption that there is a single market equilibrium has been inherited from ZIP. This means that T-ZIP agents will always set a single specific profit margin, even though there may not be an equilibrium price. Because of this, their chosen profit margin could be inappropriate in the market. Section 5.3.2 will show that, indeed, markets exist in which T-ZIP performs poorly for this reason. Section 5.4.2 provides further insight into how this occurs.

5.2 Implementation

The T-ZIP strategy described above was implemented into the T-CDA simulator (Section 4.3). In this section, the implementation is discussed by examining some example runs. The aim is to provide an illustration of market dynamics with an adaptive strategy compared to the lower bound, which is the ZI strategy. Some qualitative observations are made, which are supported by empirical results in Section 5.3.

In more detail, the example runs all have the same endowment and POS settings. Here, all sellers have a POS of 0.85 and buyers have a POS of 1. Endowments are generated from $[1.0, 1.5]$ for both buyers and sellers. Runs typically consist of 10 days. When we consider market shocks, a run consists of 20 days, with a market shock on day 11. The parameters inherent to ZIP are set

[2] Other approaches are possible, e.g. this could be the agent with the lowest POS or the median POS. However, the highest POS was chosen initially as an educated guess and appears to work well.

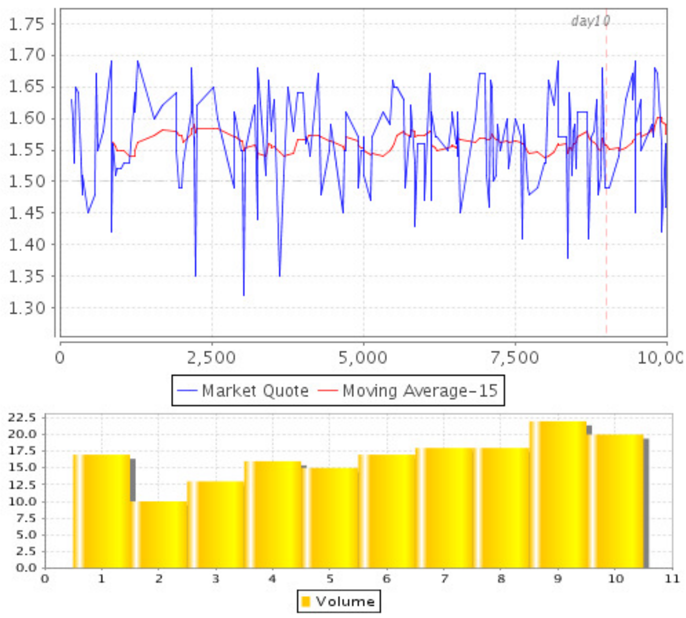

Figure 5.1: Behaviour with ZI-C.

to values randomly selected from the ranges in Table 2.1. The T-ZIP specific parameters are set to $\kappa_{\min} = 1.05$ (corresponding to $\mu = -0.05$) and $\kappa_{\max} = 2$ ($\mu = -0.5$). All these values are educated guesses suitable for a wide range of market settings and have been derived by experimental trial-and-error.

In more detail, figure 5.1 shows a run with the ZI-C strategy (Section 2.3.1). Compare this to the T-ZIP strategy (Figure 5.2 without, and Figure 5.3 with shout based bargaining). In addition to this, Figure 5.4 shows a T-ZIP run where a market shock takes place: all limit prices are shifted by one price unit. Note that the scale on these figures varies.

From these figures, we observe that the trading process is faster with the T-ZIP strategy. Furthermore, T-ZIP (with shout-based bargaining) performs the most transactions (averaging around 41 transactions per day), followed by T-ZIP without shout-based bargaining. T-ZIP agents outperform ZI, because T-ZIP buyers are not prone to overbidding: they do not bid a value they are not truly willing to pay. Moreover, T-ZIP with shout-based bargaining performs more transactions than without, because without shout-based bargaining, the market can stagnate when no commitments are being made. Also note that trading prices for T-ZIP are lower (around 1.4 instead of 1.55). This is to be expected, since sellers fail and hence it is appropriate to bid a lower value.

One important property of a learning strategy is that it can recover from market shocks. Figure 5.4 shows such a scenario. As can be seen, prices quickly converge to the new equilibrium when there is a market shock.

In a static market, we want the transaction price to converge, that is, we want to reduce the volatility of transaction prices, as is typically the case in the traditional CDA. Figure 5.8 shows that this is indeed the case. Note that here, we do not consider convergence within a trading day, but across trading days.

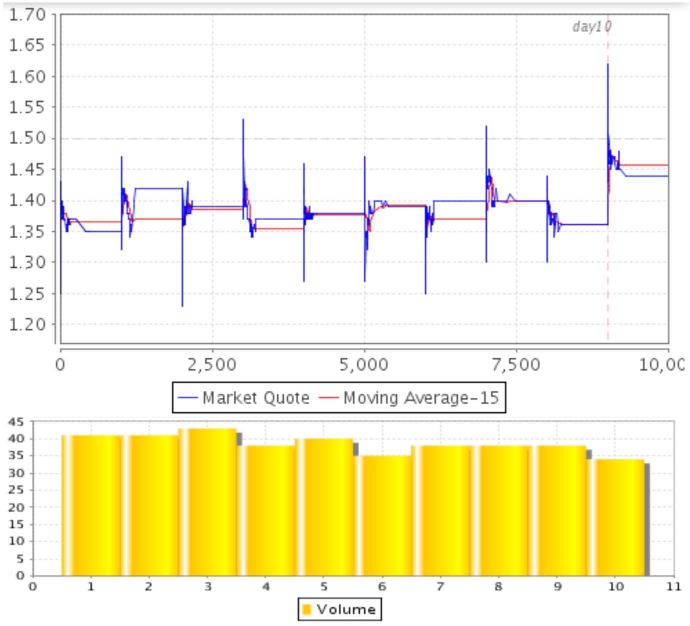

Figure 5.2: Behaviour with T-ZIP, no shout-based bargaining.

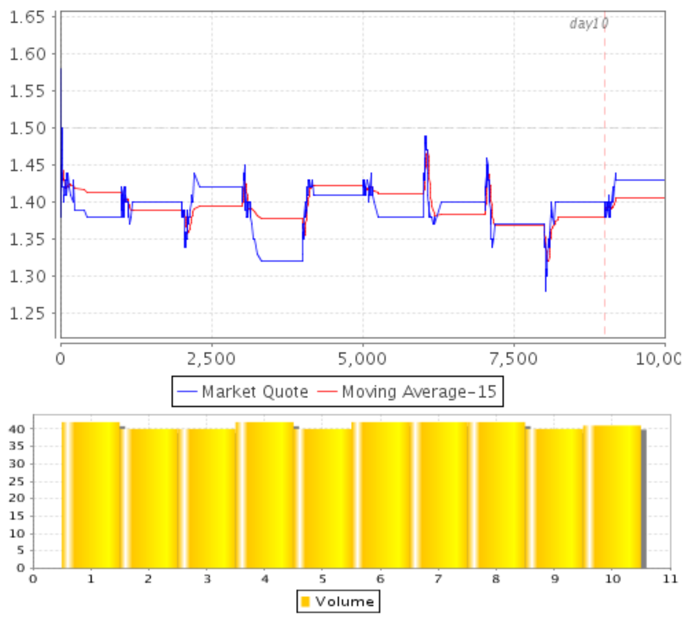

Figure 5.3: Behaviour with T-ZIP, with shout-based bargaining.

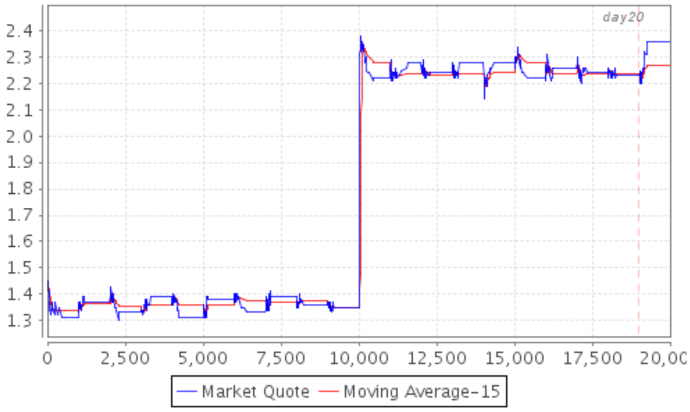

Figure 5.4: Behaviour of T-ZIP traders in event of a market shock.

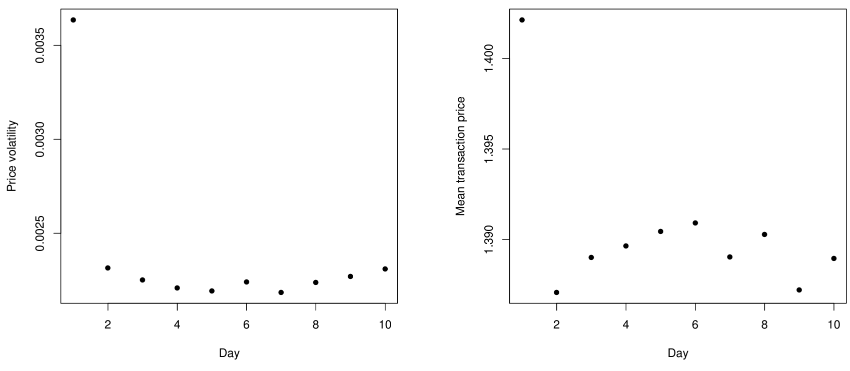

(a) Price volatility (b) Mean transaction price

Figure 5.5: Price volatility and mean transaction prices in a market with 50/50 T-ZIP traders, calculated over 100 runs.

The ZI-C strategy will typically (in a symmetric market) show convergence within a trading day (Gode and Sunder, 1993), but not between trading days.

5.3 Empirical Evaluation

Previously, we describe the design and qualitative behaviour of the T-ZIP strategy in a typical market. Now, a more complete empirical evaluation is performed. This includes:

- A comparison to the original results on the ZIP strategy (Cliff and Bruten, 1997) within a CDA-equivalent setting of our mechanism;
- A comparison of performance of T-ZIP against results with ZI agents in the T-CDA;
- An experiment where a market shock occurs;
- An evaluation of the effect of disabling the *Instantaneous Response Assumption* (recall the discussion in Section 4.3).

It is not only the T-ZIP strategy that we evaluate here. The results also provide further insights into the T-CDA mechanism and confirm that previously shown weaknesses are not due to the mechanism per se, but can be alleviated by a minimally intelligent strategy.

First, in Section 3.2, it was shown that, in general, there is no equilibrium price. Therefore, Smith's measure of price convergence (Equation 2.1) cannot be used. Hence, an alternative measure is required. Because the equilibrium price q^* does not exist, any measure of convergence can only be descriptive, whereas Smith's measure is normative. In order to measure price volatility, q^* can be replaced by the mean transaction price $\bar{q}$:

$$\tilde{\alpha} = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\tilde{q}_i - \bar{q})^2}}{\bar{q}} \tag{5.32}$$

where the $\tilde{q}_i$ are a history of n transaction prices. Because there is no normative standard against which to evaluate the transaction prices, this measure of price volatility is necessarily less informative than Smith's.

5.3.1 Comparison to ZIP

This section qualitatively compares the behaviour of the T-ZIP strategy in the T-CDA to the results in Cliff and Bruten (1997). In particular, it is shown that prices still converge to the theoretical equilibrium, in spite of the differences between the CDA and T-CDA and between the ZIP and T-ZIP strategies. Hence, the strategies are weakly equivalent (i.e. they exhibit approximately the same behaviour, although exact quantitative correspondence is not expected or required) when there is no Execution Uncertainty in the system.

In more detail, Cliff and Bruten specify four different demand and supply schedules and show that over trading days, the mean transaction price with ZIP traders converges towards the theoretical competitive market equilibrium, whereas ZI traders only converge towards this theoretical equilibrium in one of four cases. In every case, the theoretical equilibrium price is 200. The four markets are:

Symmetric demand/supply: 11 buyers and 11 sellers, with limit prices ranging from 75 to 325, in steps of 25.

Flat supply: Demand as in the previous case, but all 11 sellers have limit price 200.

Excess demand: There are 11 buyers with limit price 200 and 6 sellers with limit price 50.

Excess supply: There are 6 buyers with limit price 320 and 11 sellers with limit price 200.

The results are shown in Figure 5.6.[3] Although the standard deviation of transaction prices appears to be less for the excess demand and excess supply cases than what Cliff and Bruten report, the observed convergence is very similar to their results for all four markets.

In addition to these static markets, Cliff and Bruten investigate dynamic markets. The dynamic market they consider is one in which a market shock occurs on day 11. They consider both an increase and decrease in buyers' limit prices. Here, the experiment where the buyers' limit prices are increased is repeated. Figure 5.7 shows the result. After the shock, prices converge quite quickly, as is the case in Cliff and Bruten (1997).

In conclusion, qualitatively, the T-ZIP behaves in the same way as the ZIP, in these markets. Therefore, T-ZIP appears to be an appropriate adaptation of the ZIP to the T-CDA.

5.3.2 Comparison to ZI-C

Now, the T-ZIP strategy is compared to the ZI-C strategy, primarily in terms of market efficiency, in a wide variety of market conditions. It is shown that the T-ZIP strategy indeed alleviates the shortcoming of the T-CDA with ZI traders identified in Section 4.4.4: market efficiency does not decrease significantly when $Var(pos) = 0$ and $E(pos)$ decreases. However, it turns out that the T-ZIP strategy is not robust to high $Var(pos)$ and in some cases may even be outperformed by ZI.

Moreover, in Section 4.4, it was assumed that the seller-side failure scenario is representative of the general case. In order to support this assumption, a more complete coverage of the POS space is achieved in this experiment. Specifically, not only seller-side failure, but also buyer-side and two-sided failure are considered.

All experiments in this section were run with the simulation parameters set as for the experiments in Section 4.4 (Table 4.1). The T-ZIP parameters were the defaults (Table 2.1; $\kappa_{\text{min}} = 1.05$, $\kappa_{\text{max}} = 2$). For every condition, 20 markets were generated, consisting of a limit price and POS for every agent. Every run consists of 10 trading days. For T-ZIP, every market was repeated 500 times; for ZI, 100 times. The time limit for a trading day was set to 5000 time steps. This

[3] The format of these figures deviates somewhat from the regular format of this thesis, in order to allow direct comparison to the results by Cliff and Bruten (1997).

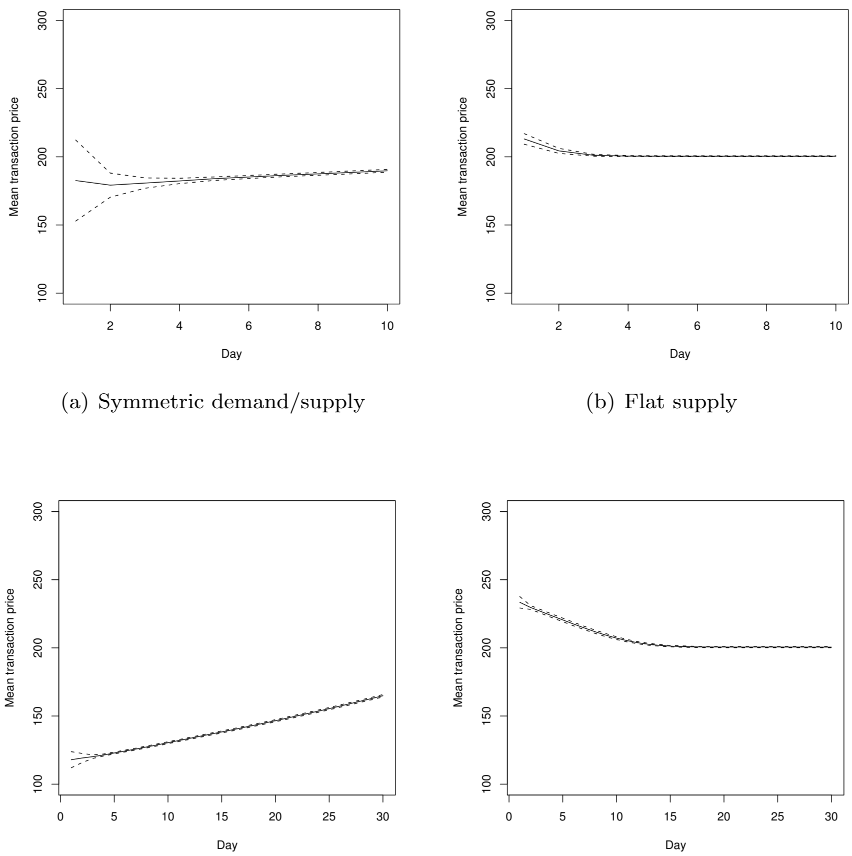

(a) Symmetric demand/supply

(b) Flat supply

(c) Excess demand

(d) Excess supply

Figure 5.6: Experiments by Cliff and Bruten repeated for T-ZIP in the T-CDA (which can be compared to Cliff and Bruten, 1997, Fig. 28–29, p. 26 and Fig 32–33, p. 27). Transaction prices should converge on the theoretical equilibrium price, 200, in all of these cases. The solid line represents the mean transaction price on a particular trading day. The dashed lines indicate the mean plus and minus one standard deviation.

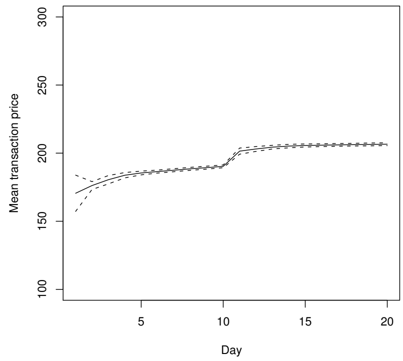

Figure 5.7: Market shock experiment by Cliff and Bruten repeated for T-ZIP (compare Cliff and Bruten, 1997, Fig. 39, p. 29). This is a symmetric demand and supply market, where every buyer limit price is increase by 50 at the start of day 11. The equilibrium price increases from 200 to 225.

means that per condition, for T-ZIP and ZI combined, a maximum of 6×10^8 time steps are required.[4]

Trust was initialised to the actual POS values, to enable meaningful comparison to the optimal allocation. Not exploring the other alternatives reduces the number of independent variables, to enable a more thorough comparison on the basis of POS settings. Specifically, both buyer-side and seller-side failure are considered, as well as two-sided failure.

One-sided failure To make this feasible, the number $Var(pos)$ settings is reduced from the experiments in Section 4.4, by choosing several settings that are distinctive in terms of both empirical results and distribution shape. Again, the Beta distribution is used. For example, Figure 4.3 shows that $Var(pos) = 0.045$ and $Var(pos) = 0.155$ produce quite dissimilar results, whereas $Var(pos) = 0.155$ and $Var(pos) = 0.205$ produce similar results. Now, $Var(pos) = 0.045$ has a bell shape and $Var(pos) = 0.155$ has a U-shape. The transition is at $Var(pos) = 5/60$, which is the uniform distribution for $E(pos) = 0.5$. See Appendix A for a discussion of the Beta distribution and the shape of the distribution for several values of $E(pos)$. Hence, the values chosen are $Var(pos) \in \{0, 0.045, 5/60, 0.155\}$.

Now, in choosing the $E(pos)$ values, $Var(pos) = 0$ must be treated separately, because at low POS ($E(pos) \leq 0.5$), no transactions are possible. Hence, the range from 0.6 to 1.0 is chosen, with steps of 0.1. For the other settings, the range 0.2 to 0.8 is chosen, with steps of 0.1. In total, there are 53 one-sided failure conditions.

[4] The original estimate of execution speed was 5×10^3 time steps per second (on the $\pm$ 2GHz AMD Opteron nodes of the University of Southampton's Iridis2 cluster). That would imply a running time of around 30 hours per condition. However, many cases run in around 15 hours whereas others may take more than 60.

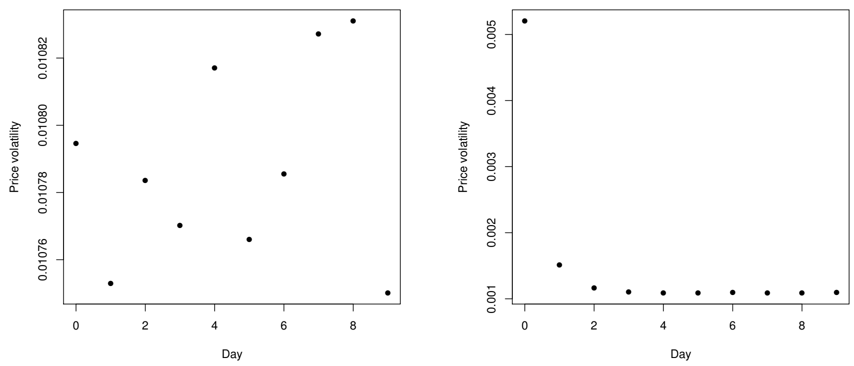

(a) ZI (b) T-ZIP

Figure 5.8: Price volatility with ZI and T-ZIP traders contrasted. Note that the scale on the y-axis differs. Market: $E(pos : S) = 1$, $E(pos : B) = 0.8$, $Var(pos : B) = 0$.

Two-sided failure These experiments are run only for $Var(pos) = 0$ and for $E(pos)$ from 0.75 to 0.95 in steps of 0.05, all combinations. The number[5] of two-sided conditions is 35. Unfortunately, it is not feasible to include the $Var(pos) > 0$ cases, as the number of conditions to run would be enormous. Furthermore, the results would be difficult to visualise and analyse, since the number of independent variables would increase from three to five.

An overview of the results is provided here. Further empirical data can be found in Appendix C.

Improved efficiency

As was noted in Section 4.4.4, market efficiency with the ZI strategy decreases when $E(pos)$ decreases beyond a certain threshold, because the bidding ranges of the ZI agents become inappropriate. Therefore, as the T-ZIP strategy was designed specifically to converge on an equilibrium price (see Section 5.1), it is expected to do well when such a price exists (i.e. when $Var(pos)$ is close to zero).

This fundamental difference between ZI and T-ZIP behaviour is illustrated by Figure 5.8, which shows price volatility (Equation 5.32) for ZI and T-ZIP markets. On the first trading day, price volatility is smaller for T-ZIP than for ZI. Moreover, price volatility with T-ZIP traders decreases further over the next trading days and converges on a certain minimal level. By contrast, ZI price volatility fluctuates randomly.

Figure 5.9 reproduces the seller-side failure condition that was also evaluated in Section 4.4.3 (see Figure 4.2(b)), for $Var(pos : S) = 0$. As was noted there,

[5] Note that the total number of conditions is 88. At 30 hours per condition, a single node would require 16 weeks to run the complete experiment.

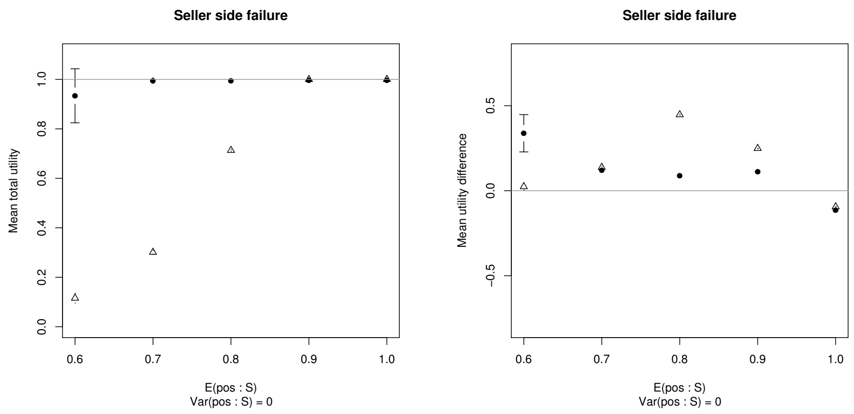

(a) Total utility (b) Utility difference

Figure 5.9: Seller-side failure, with $Var(pos) = 0$, comparing T-ZIP (circles) and ZI (triangles). Error bars indicate the 95% confidence interval, omitted where they are smaller than the plot symbol. Data was drawn from the final trading day.

ZI agents are not robust to a decreasing $E(pos : S)$. In contrast, T-ZIP is close to optimal for all levels of $E(pos : S)$.

The same pattern is visible for buyer-side failure (Figure 5.10) and two-sided failure (Figure 5.11). In each of these cases T-ZIP is robust to decreased $E(pos)$, because traders are able to adapt their shout prices to the market conditions and ensure that they do not shout a price that they would not want to transact at.

In addition to improved efficiency, we also want to achieve a fairer distribution of profits between buyers and sellers. The figures shown here are typical, with deviation from the 'zero difference' ideal usually less pronounced for T-ZIP than for ZI. However, there are exceptions to this rule. Specifically, T-ZIP deviates more when market conditions are difficult, i.e. for the extremely low values of $E(pos)$. This is to be expected, as in this case a number of agents may effectively not be realistic trading partners, creating excess demand or supply – hence, on the basis of normal market dynamics, unequal distribution of profits is to be expected. Note that T-ZIP does not exhibit the peak in utility difference for intermediate $E(pos)$ that is typical for ZI (Figure 4.3(b)). More data to support our results is provided in Appendix C.

To conclude, in Section 4.4.4, we hypothesised that the fact that the inefficiency of the market with ZI traders when $E(pos)$ is low, is caused by the inappropriate bidding behaviour of the ZI agents. Here, it was shown that a simple adaptive strategy does not exhibit the inefficiency of ZI. Therefore, the hypothesis is confirmed; this loss of efficiency is not inherent to the market and can easily be overcome by adopting a minimally intelligent strategy.

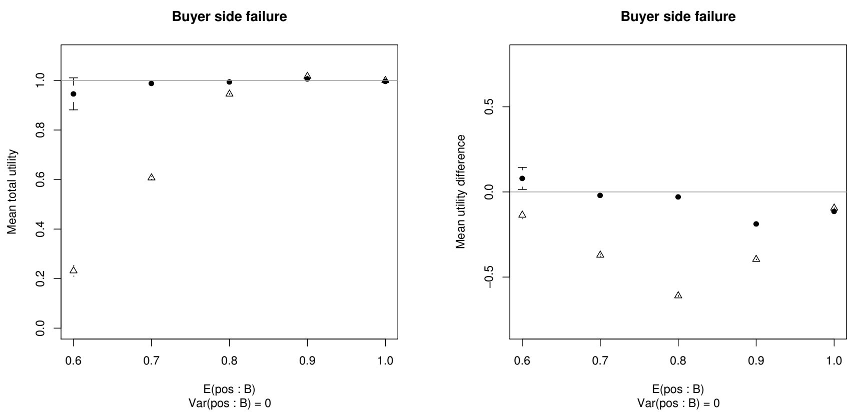

(a) Total utility (b) Utility difference

Figure 5.10: Buyer-side failure, with $Var(pos) = 0$, comparing T-ZIP (circles) and ZI (triangles). Error bars indicate the 95% confidence interval, omitted where they are smaller than the plot symbol. Data was drawn from the final trading day.

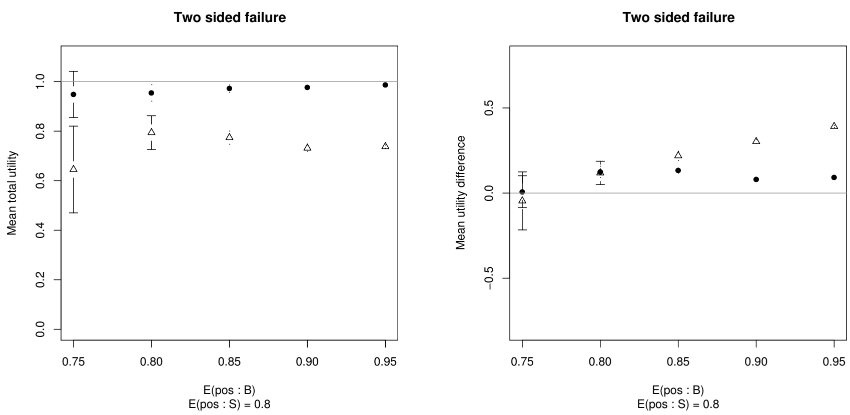

(a) Total utility (b) Utility difference

Figure 5.11: Two-sided failure, with $Var(pos) = 0$, comparing T-ZIP (circles) and ZI (triangles). Error bars indicate the 95% confidence interval, omitted where they are smaller than the plot symbol. Data was drawn from the final trading day.

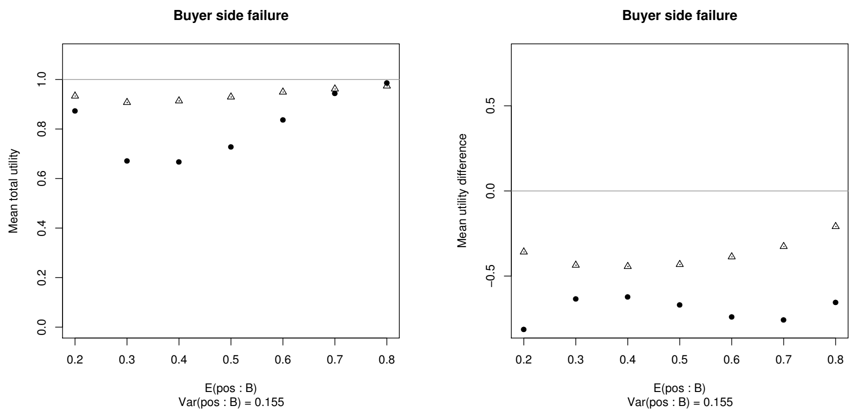

(a) Total utility (b) Utility difference

Figure 5.12: Buyer-side failure, with $Var(pos) = 0.155$, comparing T-ZIP (circles) and ZI (triangles). Error bars indicate the 95% confidence interval, omitted where they are smaller than the plot symbol. Data was drawn from the final trading day.

Failure of T-ZIP

A T-ZIP trader determines its profit margin based on the assumption that a single equilibrium (equivalent) price exists. However, when $Var(pos) > 0$, such a price does not necessarily exist. In those cases, we would expect T-ZIP to perform less well. Specifically, we look at such a case (Figure 5.12) and provide some insight into what goes wrong.

The most salient feature of the figure is the fact that ZI performs better than T-ZIP. This may seem surprising, as ZI bids randomly, whereas T-ZIP attempts to estimate a good price to bid. However, if we consider the shape of the Beta distribution when $\sigma^2 = 0.155$ (Figure A.3), it becomes apparent that in this trading environment, most buyers are clustered near the 0 and 1 POS values. Hence, the ZI-C assumption that it is safe to shout any price within the limit price is valid here (as discussed earlier, in Section 4.4.4). However, for intermediate values of $E(pos : B)$, the T-ZIP assumption that there is a single price it should converge to appears to break down. In Figure 5.13 one of these cases, $E(pos : B) = 0.5$, is shown per trading day. Indeed, transaction prices converge toward some value. However, as prices converge, total utility *decreases*. So, in this case, price convergence is harmful to the overall system.

The second important observation to be made regarding Figure 5.12 is the fact that buyer utility is consistently much higher than seller utility. This is also explained by the fact that a significant portion of the buyer population has a POS close to zero, effectively making them unavailable as trading partners. Thus, the effective buyer group is smaller than the seller group, creating excess supply. Thus, and especially with T-ZIP traders, prices are driven down towards the sellers' limit prices, creating larger profits for the buyers. Hence, in this type of market, equal distribution of utility is not possible, due to the competitive nature of markets.

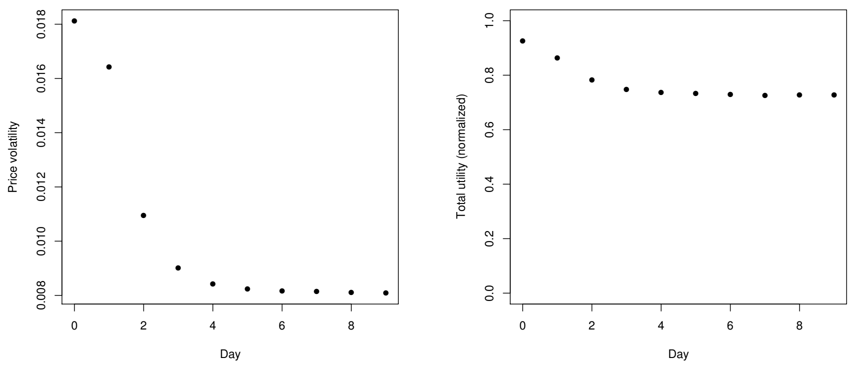

(a) Price volatility (b) Total utility (efficiency)

Figure 5.13: Harmful convergence with T-ZIP. As transaction prices converge, the market efficiency decreases. Market: $E(pos : S) = 1$, $E(pos : B) = 0.5$, $Var(pos : B) = 0.155$.

In summary, although T-ZIP clearly outperforms ZI when $Var(pos) = 0$, it is not equipped to deal with T-CDA type markets where $Var(pos) > 0$ and it may even result in lower market efficiency than the ZI strategy. This is caused by *harmful convergence*, i.e. convergence on an undesirable transaction price, as a result of which both the number of transactions occurring in the market and the market efficiency achieved decrease.

5.3.3 Instantaneous Response Step

In this section, a limited experiment is done to support the assumption that the IRS (see Section 4.3) does not impact the empirical results in any significant way. This assumption was made for two reasons:

1. In order to compare the T-CDA to the CDA (Section 4.4.3), it is important that the new mechanism with ZI traders reduces to a normal CDA given $p(a_i) = 1$; $\forall a_i \in A$. In that way, the impact of one factor (the POS) on the CDA and T-CDA was measured, without confounding the results with differences due to timing.

2. Without the IRS, experiments require a greater number of time steps to run (pilots indicate an increase of a factor 10). Therefore, statistically significant results are more readily arrived at when the IRS is enabled.

Now, it is also important to realise that the time model without the IRS is not necessarily more realistic than that with the IRS. This would depend on the scenario in which the mechanism is being evaluated. In one instance, communication delays may be significant and, hence, the additional commitment step introduces significant overhead. However, in another instance, communication delays may be much less important than the time taken to calculate potential bids, for example. If in this case the decision to accept or reject a commitment

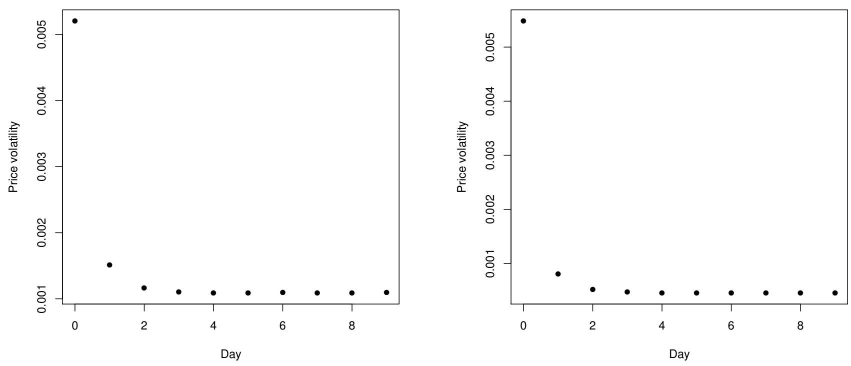

(a) IRS enabled (b) IRS disabled

Figure 5.14: Price convergence in a buyer side failure market ($E(pos : S) = 1$, $E(pos : B) = 0.8$, $Var(pos : B) = 0$). Prices still converge when the IRS is disabled.

can be made much more efficiently, the model with the IRS would be more realistic. Real world scenarios might be expected to lie somewhere in between these two extremes.

In evaluating the influence of the IRS, it is not important that the results are identical, or that the difference in some measure is statistically insignificant, but rather that the behaviour does not break down in any relevant way. Additionally, it is interesting to see whether there is a simple relationship between transaction times in the T-CDA with and without the IRS.

Because experiments without the IRS are expensive to run, even more so when ZI agents are employed, only a small number of settings is evaluated, with T-ZIP agents only. These settings are:

- No execution uncertainty (CDA);
- Buyer-side failure: $E(pos : B) = 0.8$, $Var(pos : B) = 0$;
- Seller-side failure: $E(pos : S) = 0.8$, $Var(pos : S) = 0$;
- Two-sided failure: $E(pos : S) = E(pos : B) = 0.85$.

The limited number of experiments means that results can be discussed on an individual basis.

Price volatility for IRS enabled and IRS disabled runs are shown side by side in Figure 5.14. As can be seen, convergence is preserved when the IRS is disabled. It is even true that in all four experiments, the level of price volatility reached without the IRS is slightly lower. The number of transactions and the efficiency reached are also similar. This data is available in Appendix D.

Let us now turn to the times at which transactions occur when IRS is enabled or disabled. For this, Figure 5.15 provides box plots of transaction times over 500 runs with IRS enabled and 300 runs with IRS disabled. On the left, actual transaction times are displayed. On the right, the times with IRS are multiplied

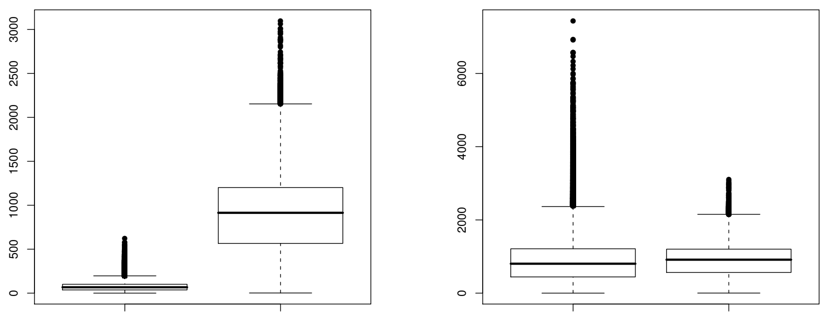

(a) Unscaled (b) Scaled

Figure 5.15: Transaction times in a buyer-side failure market. Panel (a) shows real transaction times for IRS enabled (left) and IRS disabled (right). In Panel (b), the IRS enabled transaction times are multiplied by 12.

by 12. Although the non-IRS case has a wider tail, there is a very good match between the two cases. The same holds for the other three experiments: a factor 12 aligns the boxes neatly, with the non-IRS case having a wider tail (see Appendix D for these results). Thus, it seems that there is approximately a factor 12 difference in median time to transactions, but no such relationship for the maximum time to transaction.

To conclude, far from invalidating the results obtained with the IRS enabled, these experiments suggest that the assumption is valid. If anything, T-ZIP trading behaviour is somewhat more stable with the IRS disabled. Both the lower price volatility and the smaller (relative to the interquartile range[6]) maximum transaction times provide evidence for this observation.

5.4 Discussion

The previous section details the empirical evaluation of the T-ZIP mechanism in a wide range of execution uncertainty conditions. Against this background, this section will discuss what this means for the T-CDA mechanism and why the T-ZIP strategy fails when $Var(pos) > 0$. After identifying the weaknesses of the T-ZIP strategy, approaches to building a generally applicable strategy will be proposed. Moreover, the appropriateness of the ZIP as the basis for a general T-CDA strategy will be evaluated and alternatives explored.

5.4.1 General results about the T-CDA

As was shown in Section 4.4.4, the T-CDA with ZI agents does well when the part of the agent population with which profitable trade is possible has

[6] The interquartile range is the difference between the third and first quartiles. It is a robust measure of statistical dispersion.

$E(pos) \approx 1$. This means that, when either there is no variance and $E(pos)$ is high, or when there is high variance so that most agents have a POS close to the extreme values, the T-CDA with ZI agents has a high market efficiency. For other cases, however, performance is far below the optimum. This is shown clearly by the results in Figure 5.9. It was hypothesised that this is not caused by a shortcoming of the mechanism, but rather by the fact that the ZI bidding behaviour, although appropriate in the CDA, is not appropriate when execution uncertainty is introduced. The intuition behind this is that the range of prices that are acceptable to a trader are determined not only by its limit price, but also by its trust in other traders. Specifically, ZI agents in the T-CDA may shout prices that they themselves are not willing to transact at, diminishing their chances of transacting.

Given that the ZIP strategy was designed to converge on the market equilibrium price and that the T-ZIP strategy carries this behaviour over to the T-CDA, it is expected to shout appropriate values when an equilibrium price exists. Therefore, unlike the ZI strategy, it should exhibit near-optimal efficiency when $Var(pos) = 0$. Indeed, this was shown to be the case. Hence, we may conclude that the inefficiency of the T-CDA with ZI agents found in Section 4.4.4 is due to the inappropriate bidding range of ZI agents.

Furthermore, when $Var(pos) = 0$, buyer and seller profits with T-ZIP agents are close to equality. This contrasts with ZI, where, especially for the lowest values of $E(pos)$ shown in Figure 5.9, profits tend to go mainly to the side on which failure occurs. This was also explained by the inappropriate bidding range of ZI agents. For example, when buyers fail, sellers will not be willing to trade at their cost price, but rather at a certain higher price. However, being ZI traders, they will still submit prices between their limit price and $q_{\max}$. Hence, depending on $E(pos : B)$, a certain proportion of sellers will have shouted a price that is below the price that is actually acceptable to them. Thus, if they transact at all, this is likely to take place at a price close to the limit where their expected utility is zero. The fact that T-ZIP does achieve balanced utilities in these cases confirms this explanation.

With regards to the balance of buyer and seller utility, the T-ZIP results allow several further observations to be made. The first has to do with the impact of $Var(pos)$ on the market. As can be seen in Figure 5.12, where failure is on the buyer side and $Var(pos) = 0.155$, even when the T-ZIP strategy does well, the majority of utility is allocated to buyers. In Appendix C, it can be seen that the same holds for other levels of variance, to a lesser degree. With seller-side failure, sellers are allocated a larger portion of profits. Now, as was pointed out in Section 5.3.2, the high variance means that a relatively large part of the population has a POS close to zero, making them essentially unavailable as trading partners. In this way, high $Var(pos)$ creates a market with excess supply (in case of buyer-side failure) or excess demand (seller-side failure). Thus, although the limit prices were initialised to create a symmetric market, in which buyers and sellers should have approximately equal utility, execution uncertainty can transform supply and demand in such a way as to create an unbalanced market. This is important to realise when evaluating the balance of buyer and seller utilities. Although it may be desirable to distribute utility fairly, in the case of Figure 5.12, normal market dynamics dictate otherwise. And hence, to foster a healthy competitive trading environment, we must sacrifice balanced utility in this case.

The second observation regarding the balance of buyer and seller utility is that when $Var(pos) = 0$, they are approximately equal. This means that, although the solution that optimises global welfare (Section 3.3) would minimise seller utility, in the T-CDA, sellers are still able to obtain a positive utility. Indeed, buyer and seller utilities are approximately equal. Thus, for the $Var(pos) = 0$ case, the constraint that buyer and seller utilities should be equal in calculating the optimal allocation appears to be appropriate. However, as is remarked above, under specific market conditions, this may not be appropriate as the competitive nature of markets will necessarily marginalise the profits of either buyers or sellers. In those cases, it is unclear how to define the optimal allocation. It can be argued that the optimum without the constraint of equal utilities is more appropriate, since it will at least define a standard that cannot be surpassed. However, the optimum thus found may be unrealistically high.

In conclusion, it can be said that, especially when $Var(pos) > 0$, our mechanism needs to be improved methodologically: we lack both a good baseline strategy to provide a lower bound (ZI and T-ZIP may outperform one another in different situations) and a well-defined optimal allocation to provide an upper bound. Therefore, for efficient T-CDA markets, a minimally intelligent strategy that can be generally adopted is required. Both ZI and T-ZIP are inadequate in this regard, as the first breaks down when $E(pos)$ is low and the second when $Var(pos) > 0$. Moreover, our definition of the optimal allocation is inadequate when $Var(pos) > 0$. A possible way around this can be to define the allocation that optimises social welfare (without constraints) as a standard to measure efficiency against and then to evaluate performance against the specific desiderata of the envisaged application. In any case, market behaviour and strategies cannot be evaluated purely on the basis of market efficiency, since it is an ambiguous measure when execution uncertainty is taken into account.

5.4.2 Designing T-CDA strategies

Having discussed how the results in this chapter reflect on the T-ZIP strategy and the T-CDA mechanism in general, the following will provide the intuition behind the failure of the T-ZIP strategy when $Var(pos) > 0$, and propose ways to avoid it.

The T-ZIP strategy works on the premise that there is a single transaction price that it should converge towards. When $Var(pos)$ is very low, this is a reasonable assumption. However, the experimental results show that even with moderate variance, this assumption breaks down and causes significant loss of efficiency (Definition 2.3). As was noted before in Section 5.3.2, this is caused by harmful convergence: under the learning regime, prices converge to an undesirable level. At this level, a number of potential transactions are prevented. This is what causes the diminished efficiency. We now investigate what goes wrong exactly, by analysing an example market where T-ZIP fails.

Analysis of the failure of T-ZIP

The example market we will investigate is one with buyer-side failure, where $E(pos : B) = 0.4$ and $Var(pos : B) = 5/60$. This market was chosen because it is a clear example of harmful convergence (see Figure 5.16) and analysis with buyer-side failure is more straightforward due to the form of the utility functions.

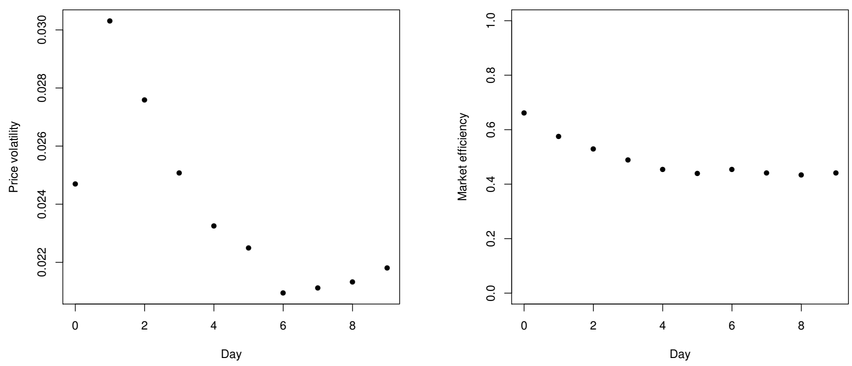

(a) Price volatility (b) Market efficiency

Figure 5.16: T-CDA market behaviour with T-ZIP traders. $E(pos : S) = 1$, $E(pos : B) = 0.4$, $Var(pos : B) = 5/60$. Market efficiency decreases as price volatility decreases. Results were calculated over 500 runs of one concrete market.

Now, as was explained before, for every condition, 20 distinct concrete markets are generated. To avoid averaging out differences between markets and thus to allow more insightful analysis, only the first such concrete market will be considered here. In this example market, Figure 5.16 shows how price volatility and market efficiency changes over trading days. As can be seen, although efficiency is not very good to start with, it decreases even further as prices converge.

In order to understand why this happens, we must find out what is causing this decrease in efficiency. It could be the case that this is caused by high transaction prices (in Section 3.3, it was shown that lower transaction prices are more efficient), or it could be the case that the number of transactions that occur is decreasing. To this end, Figure 5.17(a) is a box plot of transaction prices over trading days, while Figure 5.17(b) shows the number of transactions occurring over trading days. It appears that prices and the number of transactions are driven down. This rules out the possibility of transaction prices being too high and supports the hypothesis that the reduction in market efficiency is due to the number of transactions decreasing over trading days.

Hence, in our example the number of transactions that occur decreases over trading days, causing reduced market efficiency. We have also observed that transaction prices decrease over trading days. Now, it is vital to understand the relationship between these two observations. To support our understanding of market conditions, Figure 5.18(a) shows market structure, where buyers are represented by points on the graph, indicating their price and POS. Since sellers all have POS 1, their distribution of limit prices is too compact to be shown as individual points. Therefore, sellers' limit prices are represented by a box-and-whiskers plot. This is not very informative: we would like to see which transactions are possible in the market. Because the sellers have POS 1, buyers are willing to transact at any price below their limit price. Sellers, on the other

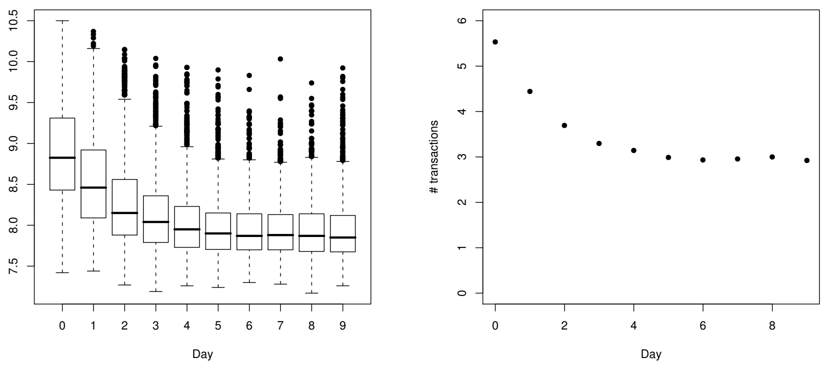

(a) Transaction prices (b) # transactions

Figure 5.17: Convergence of transaction prices and the number of transactions in the market of Figure 5.16.

hand, must trade off price and POS when deciding to transact. Using expected utility (Equation 3.2), we can find the acceptable transaction price given a seller s_j's limit price and a buyer b_i's POS:

$$\hat{u}^s_j(t_{i,j}(q)) = qp(b_i) - \ell^s_j \geq 0 \tag{5.33}$$

$$qp(b_i) \geq \ell^s_j \tag{5.34}$$

$$q \geq \frac{\ell^s_j}{p(b_i)} \tag{5.35}$$

In Figure 5.18(b), instead of the box-and-whiskers plot, the threshold price at which the seller with the lowest (at 6), the median (at 7.185), and the highest (at 7.98) limit price would be willing to transact are shown. Transactions are possible with buyers that are above and to the right of these lines. It appears that there are only 11 buyers with which a transaction with positive expected utility is possible, given the lowest limit price.

Given that we now understand the market structure well enough to predict when transactions can happen, we can analyse how convergence of transaction prices as shown in Figure 5.17(a) impacts the number of transactions. To do this, for a price q, we can draw a horizontal line that intersects the three limit lines. Then, assuming that transactions take place at this price, the intersection point determines the minimal POS at which a transaction is possible. Figure 5.19(a) shows the third quantile transaction price on day 0 (9.31) superimposed on the market structure of Figure 5.18(b). For each class of seller, transactions are only possible right of the dashed vertical lines. For the lowest limit price, the third quantile transaction price disqualifies only two potential buyers. For the median, however, all but two potential buyers are inaccessible. Contrast this with Figure 5.19(b), which shows the third quantile price for day 9 (8.12). Here, the lowest limit price allows just four transactions and the median allows none.

Given this, we understand that in Figure 5.17(b), the number of transactions decreases between trading days, because the transaction prices decrease

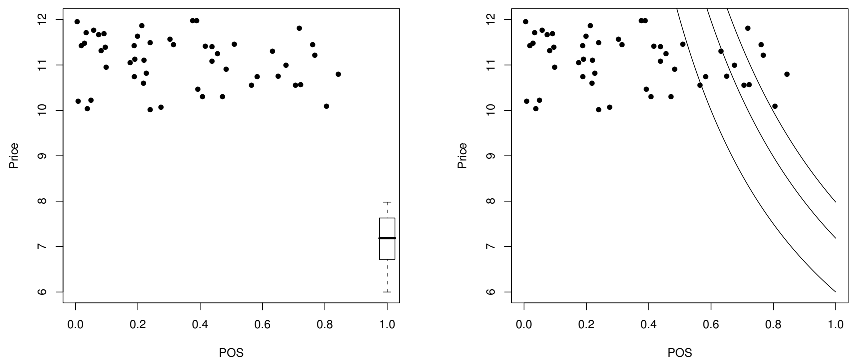

(a) Market structure (b) With seller transaction threshold

Figure 5.18: Market structure: limit prices and POS of traders in the market of Figure 5.16 visualised. Buyers are represented by circles. Sellers are represented by a box-and-whiskers plot (left) and limit lines(right). In the right-hand graph, lines (left to right) show the transaction threshold of the seller with the lowest, median and highest limit price. Transactions are possible with buyers up and right of this line.

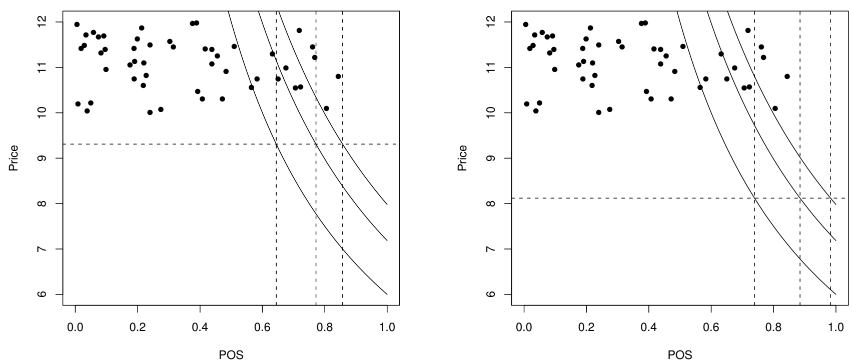

(a) Day 0 (b) Day 9

Figure 5.19: The third quantile transaction price (Figure 5.17(a)) is shown as a dashed horizontal line superimposed on Figure 5.18(b). Vertical lines are drawn where it intersects the limit curves. Only buyers to the right of these lines are accessible to the corresponding seller, at the third quantile transaction price. This shows the impact of lowering the transaction prices as happens in this market.

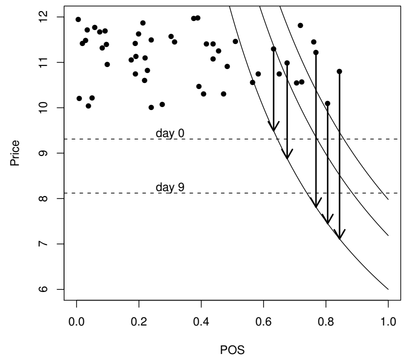

Figure 5.20: Market forces: in the market structure shown in Figure 5.18(b), excess supply lets buyers continue to decrease prices, even though this means that viable lower POS buyers lose their chances of transacting.

(Figure 5.17(a)), causing the sellers to be unwilling to transact with all but a few traders. Now, we understand how price convergence to an undesirably low level causes the number of transactions to decrease. However, it must also be explained why prices converge to this level in the first place. To understand this, we go back to the realisation that given the sellers' limit curves, there are at most 11 buyers that can transact at all. As we also noted in the previous section, this essentially creates a market with excess supply: for each of these 11 buyers, more than one seller would be willing to transact. Given this, buyers can decrease their prices until the limit curve is reached, by the argument that if one seller does not take the offer, there is another that will, because there is excess supply.

Figure 5.20 visualises that buyers drive down prices. When transaction prices are forced down, even more buyers become unavailable. Among the remaining buyers, the ones with the highest POS force the price down even further, until finally a level is reached where the number of buyers and sellers is balanced, i.e. until all sellers with a relatively high limit price are forced out of the market by the low transaction price. Thus, high POS buyers force prices down to a level that forces their lower POS competitors out of the market, even though the price these low POS buyers are willing to pay would make them attractive transaction partners to many sellers.

In conclusion, price convergence is harmful, because intra-marginal traders are forced out of the market by competitors with a higher POS. This is possible because execution uncertainty causes excess supply (or excess demand in the case of seller-side failure) and this allows buyers (sellers) to force prices down (up) to level at which some potential transactions become impossible. Note that even though buyers force prices down, as is expected in an excess supply market, this is harmful because the sellers indiscriminately learn a single price, even though transactions could take place at higher prices with buyers that have a lower POS.

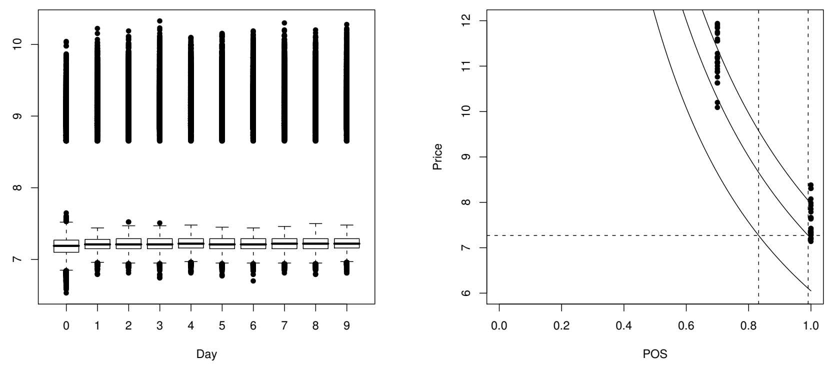

(a) Transaction prices (b) Market structure, day 0 price level

Figure 5.21: Transaction prices and market structure in a market where limit prices are highly correlated with POS. T-ZIP traders, through their convergence behaviour, again create a market with excess supply.

Failure in a segmented market

From the above discussion, it is clear that the inefficiency of T-ZIP is caused by its very nature: convergence on a single price. This in itself is not surprising and in fact this assumption was known in advance to be false. However, the way in which failure occurs is more surprising. One might expect the strategy to fail to distinguish between two sets of traders and thereby arrive at a price level that is inappropriate for both groups.

Specifically, an experiment was run to observe this failure type. In this experiment, seller limit prices are still drawn from $[6, 8]$, but buyers are separated into two groups of 25 agents each. The first group has POS 0.7 and limit prices from $[10, 12]$, and the second group has POS 1 and limit prices from $[7, 8.4]$. From the point of view of sellers, trade with either group should be desirable, but at different price levels. Now, the presence of these two groups could prevent price convergence.

Indeed, prices are more volatile ($\alpha^* \approx 0.02$) in this market than in those studied before (usually $0.001 \leq \alpha^* \leq 0.01$). However, as can be seen from a plot of transaction prices and the market structure (Figure 5.21), prices clearly converge - to the lowest price level. Hence, the group with lower POS has a diminished chance of transacting. So, even in this market, the market behaviour can be explained by the fact that an excess supply situation is created by the convergence of prices, not by a failure to converge to an appropriate price.

The intuition is as follows: several examples of trade at the lower price level will cause sellers to assume they are pricing themselves out of the market. Thus, they will gradually lower their profit margin. The margin needs only be lowered a little to exclude the low POS buyer group from trade. Thus, the more sellers lower their margin below this limit, the more examples of trade at the lower price level will occur and the more sellers will drop below this margin.

Therefore, although having two buyer groups of equal utility but different

POS does cause suboptimal behaviour, this is explained in the same fashion as for less structured markets. Indeed, prices do not fail to converge, as one might expect, but rather favor the group with higher POS.

Designing a general strategy

In case the failure is caused by prices failing to converge, a solution would be to let traders group their opposing traders and let them learn an appropriate price for each group. However, as is shown above, a continuum of different levels of POS and excess supply or excess demand may also cause the T-ZIP strategy to fail. In this case, a solution is less obvious. Perhaps one of the challenges is to ascertain whether markets in general resemble the market of Figure 5.21(b), or that of Figure 5.18(b). For now, we assume that either may occur and, specifically, we assume we may encounter markets where there is very little structure on which to base a 'grouping' of agents, as in Figure 5.18(b).

In that case, a possible solution is to make the profit margin a function of the agent's trust in the trading partner. Of course, a trader will not be able to gather experience on all different levels of trust and hence there will need to be some sort of interpolation to provide the appropriate margin for each level of trust. However, the ZIP learning rules would need to be replaced by another set of rules, based on machine learning.

Alternatively, the limit price could be made a function of the trust a trader places in the other party, letting this be the zero-utility price for each level of trust. In that case, the profit margin could be a fixed value. However, if the margin is fixed, this may lead to the agent failing to take advantage of a segment of the market where a greater margin is possible. On the other hand, it has the advantage that no safeguards are required to make sure that the profit margin does not imply a negative utility for the agent.

Finally, the two approaches can be combined by having both the profit margin and the limit price as functions of the trust value. This would bring the advantage of fine-tuning the margin for specific levels of trust and not having to safeguard against setting the margin too low. However, it might also be the most dependent on the trust model and therefore require the highest level of reliability in trust information. Specifically, the interaction of two functions that depend on trust may amplify errors present in the trust information.

In either case, the bidding strategy will need to be revised. It is no longer appropriate to just submit a bid based on the profit margin, as the price that we should bid is now a function of trust. Thus, a trader needs to take the current outstanding bids and asks into account as well as the reliability of the traders. Perhaps a technique to estimate the chances of transacting, similar to what is done by the GD strategy (Gjerstad and Dickhaut, 1998), could be useful (see Section 2.3.3). On the other hand, if a technique where the profit margin depends on trust is adopted, the trader could simply target the agent(s) that would allow the highest profit margin.

Generalising from the above discussion, whichever path of adaptation in the T-CDA is chosen, it would seem that a general trading strategy would hardly resemble ZIP any more. T-ZIP has been valuable in pointing out what is required of a trading strategy, but the value of ZIP as a basis for a general strategy is questionable. Therefore, it seems prudent to also take notice of techniques used by other strategies such as those discussed in Section 2.3.3, as well as the more general machine learning literature.

5.5 Summary

In this chapter, a simple adaptive trading strategy for the T-CDA was designed, with the aim of showing that high efficiency is achievable by the T-CDA mechanism even when $E(pos)$ is low. Results with ZI traders (Section 4.4.4) seem to suggest otherwise, but it was hypothesised that the low efficiency of the T-CDA in those cases is attributable to the simplistic and uninformed bidding behaviour of ZI agents. The T-ZIP strategy developed in this chapter is designed to converge to the appropriate transaction price and hence, it should achieve close to optimal results when $Var(pos) = 0$.

However, it is not designed with the $Var(pos) > 0$ case in mind. In more detail, it inherits the assumption that a single market equilibrium exists from the ZIP strategy. Therefore, T-ZIP could fail when this is not the case.

The T-ZIP strategy is a relatively straightforward adaptation of the ZIP trading strategy for the CDA to the T-CDA. A new way to set the target price for the learning algorithm was derived as part of this work and, because market events in the T-CDA are different from the CDA, the events that trigger the learning mechanism were also redefined.

Several example runs were examined and show a significant qualitative difference with ZI-C trading behaviour, as well as improved efficiency and a greater number of performed transactions. This impression is confirmed by an empirical evaluation. When $Var(pos) = 0$, T-ZIP indeed derives close to optimal efficiency, even when $E(pos)$ is small. Furthermore, there is clear convergence of transaction prices. However, as expected, when $Var(pos) > 0$, the T-ZIP traders are not efficient. It turns out that prices do converge, but to an undesirable value. Because of this harmful convergence, efficiency decreases between trading days.

In addition to an evaluation of the trading mechanism, the influence of the IRS is examined in a limited experiment. If anything, price volatility is lower when IRS is disabled than when it is enabled. Efficiency and the number of transactions do not differ in a relevant way. The timing of transactions was investigated and a rough factor 12 slowdown (in terms of the number of time steps required to complete a trading round) when disabling IRS was found. This should be interpreted carefully, as only a small number of cases was investigated. However, the data does suggest that the results in this thesis will still be valid when the IRS is rejected.

Finally, suggestions for the development of a generally applicable trading strategy for the T-CDA were discussed, as well as alternatives to the ZIP as the basis for derivation.

Chapter 6

Conclusions

In this thesis, the Trust-Based CDA (T-CDA) was developed. This is a novel, decentralised, trust-based market mechanism derived from the CDA. The T-CDA is robust to execution uncertainty, by allowing agents to use a trust model in their decision making. Moreover, a rudimentary strategy for the T-CDA, Trust-Based ZIP (T-ZIP), was developed. Although it is not a generally applicable strategy for the T-CDA, it provides further insights into the market and a starting point for more general strategies.

In what follows, the research contributions made in this thesis are evaluated against the research objectives (Section 1.3). Then, directions for further work are identified.

6.1 Conclusions

In this section, the research objectives stated in Section 1.3 are revisited and we discuss how this thesis addresses each of them in turn. Then, we evaluate how this reflects on the main research question.

Objective 1. The first research objective is to create a novel trading mechanism, based on the CDA, that is robust to execution uncertainty, by allowing agents to use a trust model in their decision making during the trading process.

In Chapter 4, the T-CDA mechanism was outlined. The T-CDA separates the commitment from the bidding process: in the CDA, when an agent submits a bid or ask, the agent is automatically committed to a transaction with any agent it is matched with. In the T-CDA, commitments are made explicit. To transact, in addition to submitting a bid or ask with a competitive price to the market, the trader has to indicate its willingness to transact by committing to a transaction with a specific other trader.

Thus, in the T-CDA, agents can differentiate between their potential transaction partners. This allows them to use their beliefs about the reliability of others in choosing a specific transaction. These beliefs can be modelled by a probabilistic trust function. Then, agents can balance price and execution uncertainty in their decision making, by estimating the expected utility of transacting.

In this way, the T-CDA addresses execution uncertainty in a decentralised manner: it merely provides a framework in which an agent can express its desires effectively. When it comes to dealing with uncertainty in the market, the onus is placed entirely on the agent. Specifically, the agent is responsible for making sure their trust function is accurate and to behave in a strategic way in order to derive maximum profit.

The T-CDA is empirically shown to be robust to execution uncertainty, in Section 4.4. Specifically, the CDA breaks down when execution uncertainty increases in the system: agents derive negative utility. It is shown that in the T-CDA, this does not happen: on average, agents will derive a positive payoff if this is at all possible.

Objective 2. After developing the T-CDA, the second objective is to study its properties with minimally intelligent traders. In Section 4.4, it is shown that, even with ZI traders, even when reliability differences between agents are very large, the mechanism derives a profit that is close to optimal as long as the agents with whom profitable transactions are possible have a POS close to 1. However, efficiency decreases significantly when this is not the case. It is argued that this is caused by the inappropriate bidding behaviour of ZI agents. In particular, buyers bid prices that are too high and sellers ask prices that are too low. At these prices, they themselves are not willing to transact, causing decreased overall utility. Thus, the ZI strategy is not an adequate benchmark of the T-CDA.

This is addressed in Section 5.3, where it is shown that the decrease of efficiency shown with ZI traders is avoided by the T-ZIP traders developed in Chapter 5. Specifically, with T-ZIP traders, efficiency is close to the optimum, whenever $Var(pos)$ is close to zero. Hence, social welfare is close to optimal in the T-CDA, given reliable trust information and given an adequate trading strategy.

Another desideratum is individual rationality, which is addressed in the T-CDA by giving the agents control over who they transact with. Thus, agents can ensure they only transact when they expect a positive utility. Since the system achieves positive social welfare and does not discriminate between agents in any way, if an agent has gathered reliable trust information, it can expect a positive payoff from participating in the mechanism.

Related to this is the balance of utilities derived by buyers and sellers. In general, if there is balanced (symmetric) demand and supply, and both buyers and sellers employ adequate strategies, they will derive approximately equal utilities (see Section 5.3). However, as is pointed out there, the POS distribution transforms the demand and supply and hence some of the investigated markets have excess demand or excess supply – severely skewing the balance of profits. This is unavoidable in a market, as it is a direct result of the competitive nature of such institutions.

Finally, in Section 4.4.5, it is shown that the mechanism is relatively robust to errors in the trust information. That is, performance degrades linearly with the error that is introduced on the trust function.

Objective 3. Complementary to the new mechanism developed in this thesis, is a trading strategy that operates within the new mechanism. In Chapter 5, the T-ZIP strategy is detailed. It is an adaptation of the ZIP strategy for the traditional CDA. Designing the T-ZIP was a significant challenge, as the opportunities for learning presented by trade in the CDA do not translate directly to the T-CDA. In addition to this, in the T-CDA, prices cannot be taken at face value and an agent must therefore calculate an equivalent price in order to adjust its own profit margin based on prices in the market.

Objective 4. Finally, we aim to study the properties of T-ZIP and to benchmark its performance. The T-ZIP strategy is shown to exhibit the same behaviour as the original ZIP strategy (Section 5.3.1). That is, price volatility rapidly decreases during trade and the strategy recovers after a market shock. Furthermore, even in non-symmetric markets, prices eventually converge on the theoretical equilibrium price.

As was mentioned above, the T-ZIP achieves close to optimal efficiency when $Var(pos)$ is close to zero. However, it is also shown to break down when $Var(pos)$ is large (Section 5.3.2). This is a direct result of the price convergence behaviour the T-ZIP strategy exhibits. Specifically, the price a T-ZIP agent shouts is indiscriminate to the current population of the market. Therefore, as prices converge, an increasing number of traders is excluded from taking part in transactions.

When taken together, we have seen how this thesis meets its stated research objectives. Now, we will see how the overall research question is answered. In Section 1.3, the research question is formulated as follows: may the Continuous Double Auction be extended by a component that enables agents to balance cost and reliability of the transactions they agree to, by incorporating a trust-model in their decision making, in a way that achieves close to optimal social welfare even when faced with execution uncertainty?

Based on the work in this thesis, the answer to this question is yes. The T-CDA is such an extension of the CDA. However, the potential of the T-CDA is not yet realised completely, for lack of a generally applicable trading strategy. Thus, further work is required to give a more definitive answer to the stated research question. However, we believe the design of an adequate strategy for the T-CDA is feasible.

6.2 Further Work

The previous section outlines the contributions made by this thesis. Now, we outline directions for further work. First and foremost, as is pointed out above, the T-CDA currently lacks an adequate strategy that realises the full potential of the mechanism and that can serve as a benchmark to evaluate other, possibly more advanced, trading strategies against. The development of a generally applicable trading strategy should be the top priority, since without it market efficiency may, in specific conditions, reduce to a fraction of the optimum. Hence, currently there is little incentive to use the T-CDA in a real-world application.

Second, the T-CDA has thus far been evaluated using a simulated trust model. One of the possible sources of trust information is market events. Not only can agents learn about the reliability of others by direct interactions, but the commitment behaviour of others can be observed and possibly used as an indirect source of reputation: unreliable agents will be shunned by knowledgeable traders. Alternatively, a more elaborate simulation of trust models could be implemented, showing how certain properties of the trust model impact on the mechanism. For example, the trust model could have a certain rate at which it converges toward the true POS of an agent. Then, it is interesting to investigate the effect of a sudden or gradual change in agents' POS on the market.

Third, in order to be able to compare the results obtained with the T-CDA against an optimal solution, the model in which the T-CDA is evaluated is currently simplified in various ways. Since the T-CDA was contrived to be a resource allocation mechanism for large scale, open, distributed systems, it should be evaluated in more realistic scenarios where some, or all, of these assumptions are relaxed. In such a context, the T-CDA should be compared to other resource allocation mechanisms.

Appendix A

Beta Distribution

The beta distribution is a family of continuous probability distributions defined on $[0, 1]$. It has two shape parameters, α and β, which must be positive. Depending on these parameters, the distribution takes on different shapes (Table A.1).

Now, if we know the distribution mean μ and variance σ^2, the parameters α and β can be calculated, as was shown in Teacy (2006, app. B):

$$\alpha = \frac{\mu^2 - \mu^3}{\sigma^2} - \mu \tag{A.1}$$

$$\beta = \frac{\alpha}{\mu} - \alpha \tag{A.2}$$

Since the parameters must be positive, with Equation A.1 and $\alpha > 0$ we can bound the variance σ^2 that can be achieved with a beta distribution for a specific μ:

$$\frac{\mu^2 - \mu^3}{\sigma^2} - \mu > 0 \tag{A.3}$$

$$\mu^2 - \mu^3 > \mu\sigma^2 \tag{A.4}$$

$$\mu - \mu^2 > \sigma^2 \tag{A.5}$$

Note that the $\beta > 0$ constraint will not introduce any additional constraints on μ and σ^2.

Having introduced the beta distributions and how to find its parameters α and β from the desired μ and σ^2, the distribution shapes for three important settings of σ^2 are illustrated on the following pages.

<table>
<tr><td>$\alpha = 1$</td><td>$\beta = 1$</td><td>uniform</td></tr>
<tr><td>$\alpha < 1$</td><td>$\beta < 1$</td><td>U shape</td></tr>
<tr><td>$\alpha < 1$</td><td>$\beta \geq 1$</td><td rowspan="2">strictly decreasing</td></tr>
<tr><td>$\alpha = 1$</td><td>$\beta > 1$</td></tr>
<tr><td>$\alpha = 1$</td><td>$\beta < 1$</td><td rowspan="2">strictly increasing</td></tr>
<tr><td>$\alpha > 1$</td><td>$\beta \leq 1$</td></tr>
<tr><td>$\alpha > 1$</td><td>$\beta > 1$</td><td>unimodal (bell shape)</td></tr>
</table>

Table A.1: Beta distribution shape depends on α and β.

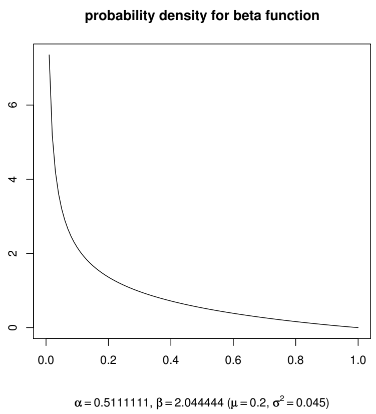

(a) $\mu = 0.2$

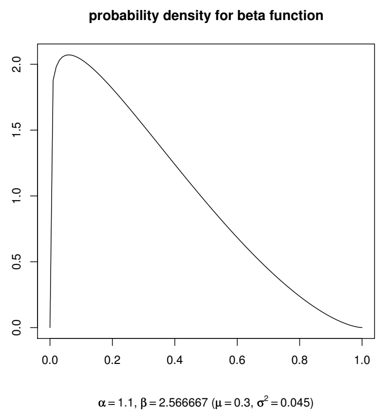

(b) $\mu = 0.3$

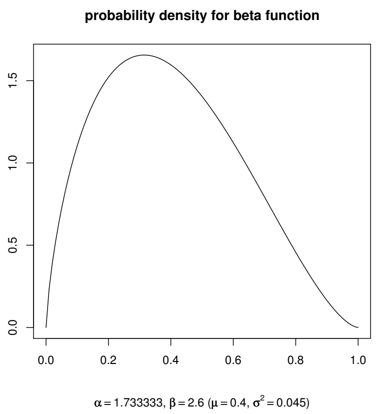

(c) $\mu = 0.4$

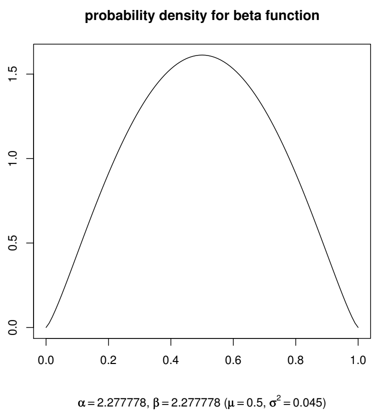

(d) $\mu = 0.5$

Figure A.1: Beta distribution for $\sigma^2 = 0.045$. Plots for $\mu > 0.5$ are omitted because of symmetry.

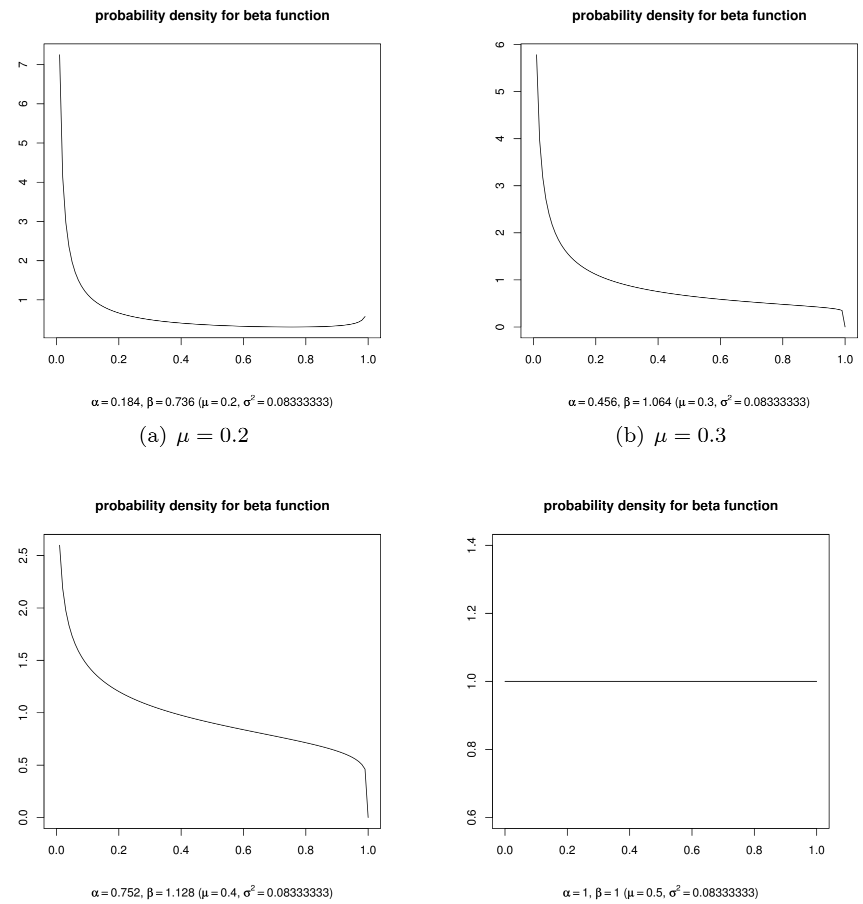

(a) $\mu = 0.2$

(b) $\mu = 0.3$

(c) $\mu = 0.4$

(d) $\mu = 0.5$

Figure A.2: Beta distribution for $\sigma^2 = 5/60$. Plots for $\mu > 0.5$ are omitted because of symmetry.

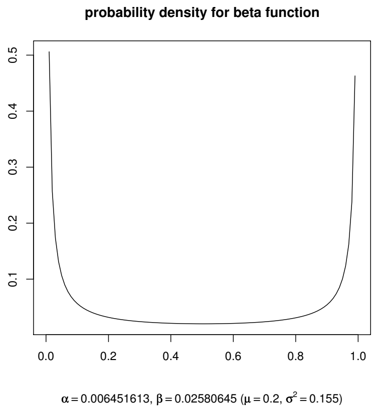

(a) $\mu = 0.2$

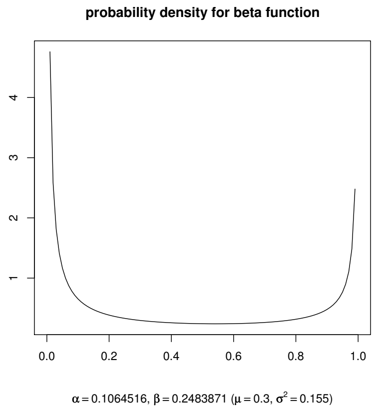

(b) $\mu = 0.3$

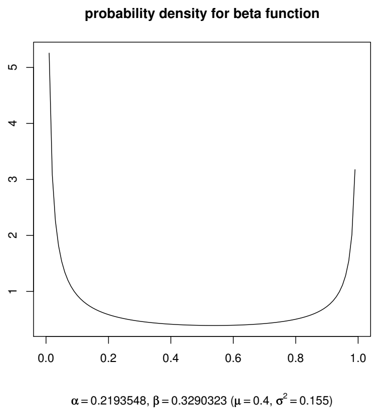

(c) $\mu = 0.4$

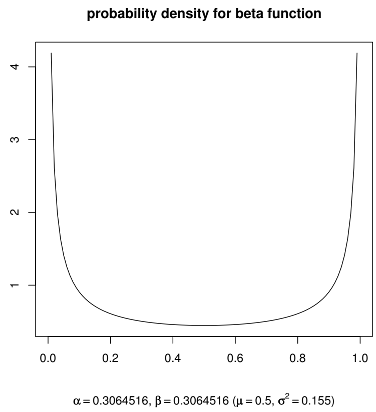

(d) $\mu = 0.5$

Figure A.3: Beta distribution for $\sigma^2 = 0.155$. Plots for $\mu > 0.5$ are omitted because of symmetry.

Appendix B

T-CDA: CDA Comparison

This appendix overviews experimental data that goes with the comparison of the T-CDA with the traditional CDA, described in Section 4.4.3. Each of the tables in this appendix compares the PERFECT trust condition to another condition on realised buyer utility. Comparisons were run at different POS levels (rows) and different levels of variance (columns).

For every comparison, two t-tests were done. In either case, the null hypothesis was that the means are equal. For one test, the alternative is that PERFECT has a higher mean, for the second that PERFECT has a lower mean. The 'c' column gives results at the $\alpha = 0.05$ level, which would be appropriate if we run just one comparison. The 'e' column, on the other hand, gives results at the $\alpha' = 1 - 0.95^{1/65}$, which protects the null hypothesis against spurious results. A '+' indicates that PERFECT has a significantly higher mean, '-' indicates a lower mean and '0' indicates no significant difference.

In general, given the number of comparisons, spurious results are expected to occur at the α level, especially in the high POS cases, where we expect little or no difference because the problem is 'easy' in these cases: transactions are generally desirable. In order to provide the most informative results, in stead of providing only the level α' or α results, both are included, allowing the differences to be analysed.

Specifically, it must be noted that there are two instances in Table B.2 (comparison of PERFECT and NAIVE) where NAIVE has a significantly higher mean buyer utility at the α level. However, these results are insignificant even if we view the eight instances that show no significance at the α' level as a separate experiment in which we try to show that naive does better. That is, these results are insignificant even at $\alpha^* = 1 - 0.95^{1/8}$. Also note that both occur for high values of POS and that therefore, true means are likely to be equal. Thus, these results are likely to be spurious. The same discussion applies to similar cases in other tables.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10	+	+	+	+	+	+	+	+												
0.30	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.50	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0.60	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
0.70	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.75	+	+	+	+	+	+	+	+	+	+	+	+	+	+						
0.80	+	+	+	+	+	+	+	+	+	+	+	+								
0.85	+	0	+	+	+	+	+	+	+	+										
0.90	+	+	+	+	+	+	+	+												
0.95	+	+	+	+	+	+														
1.00	+	+																		

Table B.1: Comparison of PERFECT and RANDOM trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10	+	+	+	+	+	+	+	+												
0.30	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.50	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0.60	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
0.70	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.75	+	+	+	+	+	+	+	+	+	+	+	+	+	+						
0.80	+	0	+	+	+	+	+	+	+	+	+	+								
0.85	-	0	+	0	+	+	+	+	+	+										
0.90	0	0	0	0	+	+	+	+												
0.95	-	0	0	0	+	+														
1.00	0	0																		

Table B.2: Comparison of PERFECT and NAIVE trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10					0	0	+	+												
0.30					+	0	0	0	+	+	+	+	+	+	+	+				
0.50			0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0.60	0	0	+	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
0.70	0	0	0	0	+	+	+	+	+	+	+	+	+	+	+	+				
0.75	0	0	+	0	+	+	+	+	+	+	+	+	+	+						
0.80	0	0	+	+	+	+	+	+	+	+	+	+								
0.85	0	0	+	+	+	+	+	+	+	+										
0.90	0	0	0	0	+	+	+	+												
0.95	-	0	0	0	+	+														
1.00	0	0																		

Table B.3: Comparison of PERFECT and MEAN trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10					0	0	0	0												
0.30					+	0	0	0	0	0	0	0	0	0	0	0				
0.50	+	0	+	+	+	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.60	+	0	+	0	+	0	0	0	+	0	+	0	0	0	+	0	0	0		
0.70	+	+	+	+	+	0	+	0	0	0	0	0	0	0	0	0				
0.75	+	+	+	+	+	0	+	0	0	0	0	0	0	0						
0.80	+	+	+	+	0	0	+	+	0	0	0	0								
0.85	0	0	0	0	0	0	0	0	0	0										
0.90	0	0	0	0	0	0	0	0												
0.95	0	0	0	0	0	0														
1.00	0	0																		

Table B.4: Comparison of PERFECT and NOISE05 trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10	+	+	+	+	+	+	0	0												
0.30	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0				
0.50	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0	0	0	0
0.60	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0	0		
0.70	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0	0				
0.75	+	+	+	+	+	+	+	+	+	+	0	0	0	0						
0.80	+	+	+	+	+	+	+	+	0	0	0	0								
0.85	-	0	+	0	+	+	+	+	0	0										
0.90	0	0	0	0	+	0	0	0												
0.95	0	0	0	0	0	0														
1.00	0	0																		

Table B.5: Comparison of PERFECT and NOISE15 trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10	+	+	+	+	+	+	+	+												
0.30	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.50	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0.60	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
0.70	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.75	+	+	+	+	+	+	+	+	+	+	+	+	+	0						
0.80	+	+	+	+	+	+	+	+	+	+	+	+								
0.85	-	0	+	0	+	+	+	+	+	0										
0.90	0	0	0	0	+	+	0	0												
0.95	0	0	0	0	0	0														
1.00	+	0																		

Table B.6: Comparison of PERFECT and NOISE25 trust conditions on realised buyer utility.

	0.000		0.010		0.045		0.085		0.125		0.155		0.185		0.205		0.235		0.245	
	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e	c	e
0.10	+	+	+	+	+	+	+	+												
0.30	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.50	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0.60	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
0.70	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+				
0.75	+	+	+	+	+	+	+	+	+	+	+	+	+	+						
0.80	+	+	+	+	+	+	+	+	+	+	+	+								
0.85	-	0	+	+	+	+	+	+	+	+										
0.90	0	0	0	0	+	+	+	+												
0.95	0	0	0	0	+	+														
1.00	+	0																		

Table B.7: Comparison of PERFECT and NOISE40 trust conditions on realised buyer utility.

Appendix C

T-ZIP: ZI-C Comparison

This appendix provides a more complete overview of the data obtained in the experiment comparing the T-ZIP strategy to ZI-C (see Section 5.3.2). Although many hundreds of pages could be dedicated to plots of price volatility over trading days, distribution of transaction times, distribution of transaction prices, convergence of efficiency or the number of transactions (and indeed this was done as part of the analysis of this experiment), only the essential data is provided here.

Specifically, for the one-sided failure experiments, plots of last-day buyer and seller efficiency, as well as overall efficiency and discrepancy of buyer and seller utilities are provided. For the two-sided failure experiment, only plots of overall efficiency are provided, as T-ZIP will achieve near-perfect balance of utilities in every case, whereas ZI traders exhibit an imbalance as expected from the one-sided experiments (that is, the lower POS group extracts more profits).

In all plots, error bars indicate a 95% confidence interval and are omitted where they would be smaller than the plot symbol.

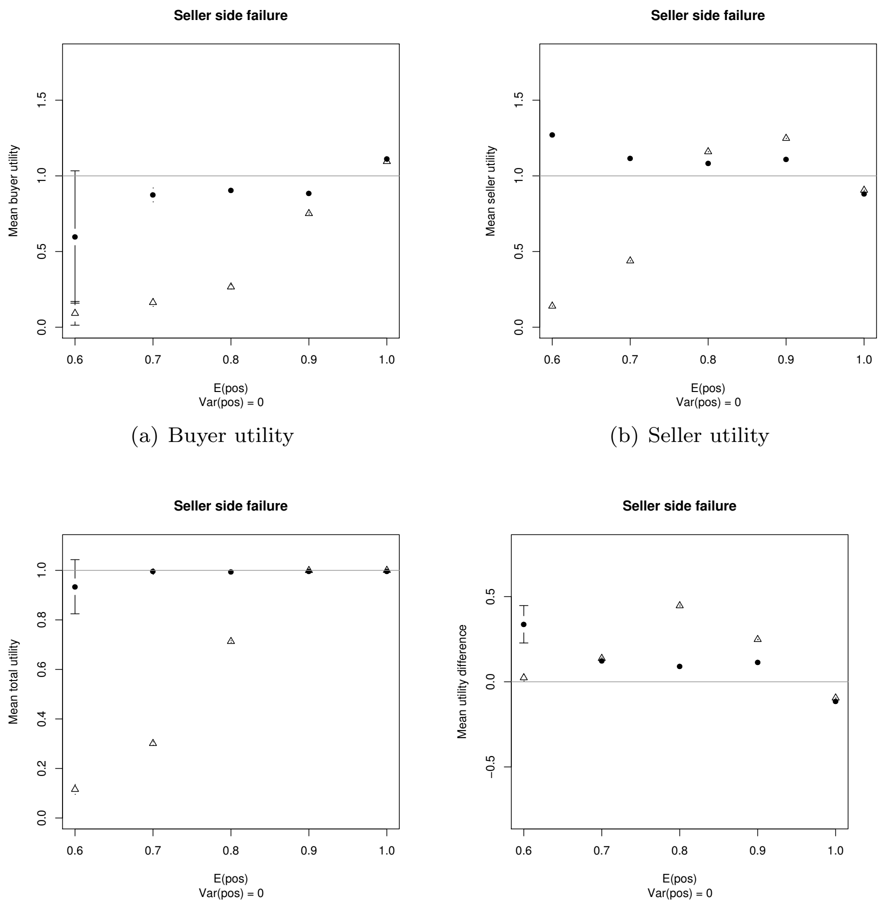

(a) Buyer utility

(b) Seller utility

(c) Total utility

(d) Utility difference

Figure C.1: Seller-side failure, with $Var(pos) = 0$, comparing T-ZIP (circles) and ZI (triangles).

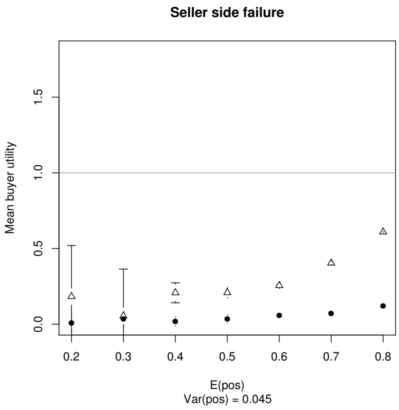

(a) Buyer utility

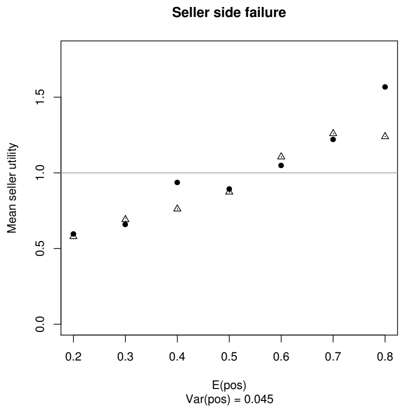

(b) Seller utility

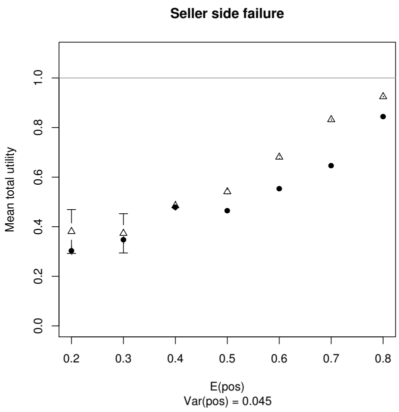

(c) Total utility

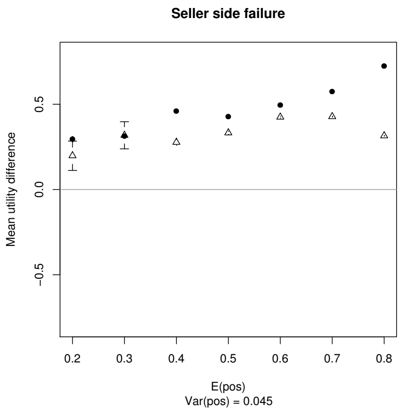

(d) Utility difference

Figure C.2: Seller-side failure, with $Var(pos) = 0.045$, comparing T-ZIP (circles) and ZI (triangles).

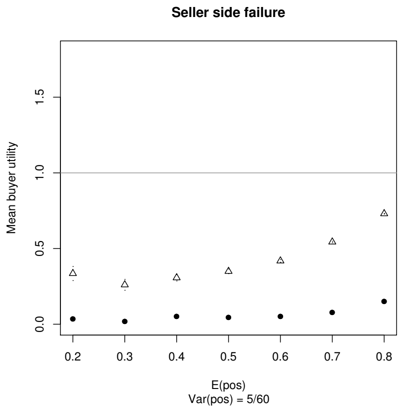

(a) Buyer utility

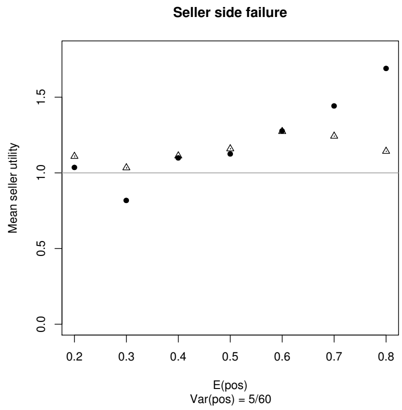

(b) Seller utility

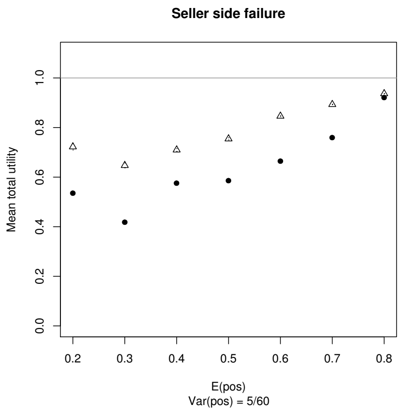

(c) Total utility

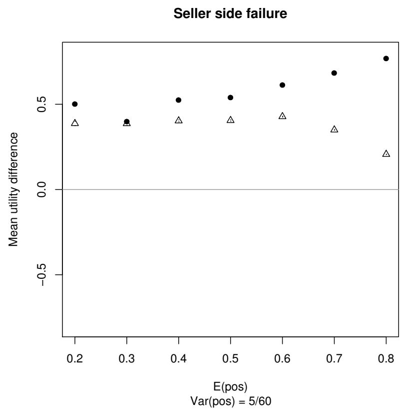

(d) Utility difference

Figure C.3: Seller-side failure, with $Var(pos) = 0.083$, comparing T-ZIP (circles) and ZI (triangles).

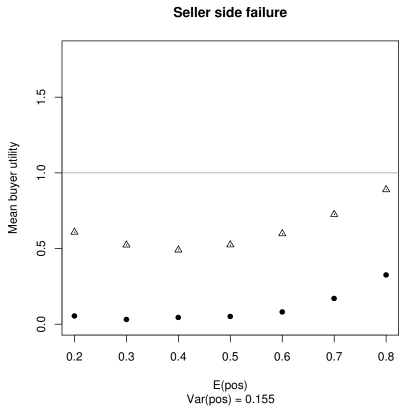

(a) Buyer utility

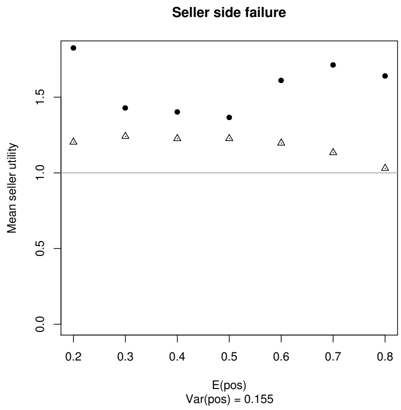

(b) Seller utility

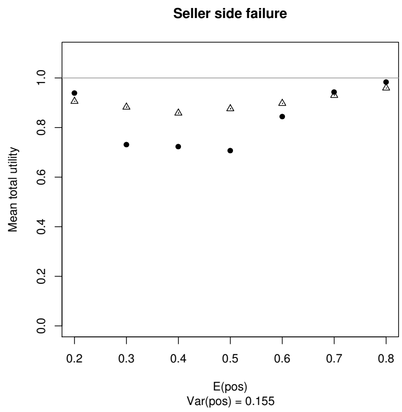

(c) Total utility

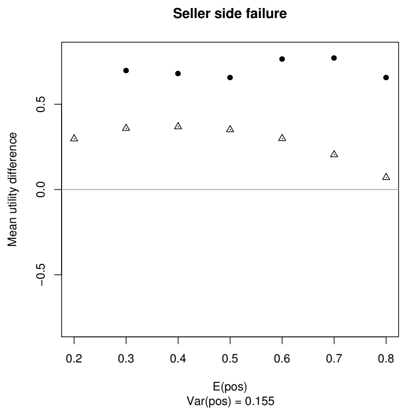

(d) Utility difference

Figure C.4: Seller-side failure, with $Var(pos) = 0.155$, comparing T-ZIP (circles) and ZI (triangles).

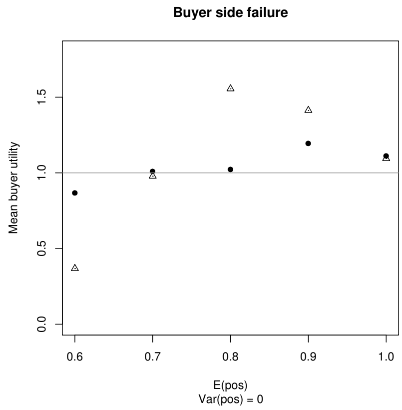

(a) Buyer utility

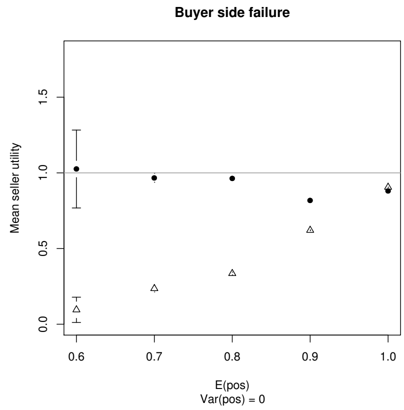

(b) Seller utility

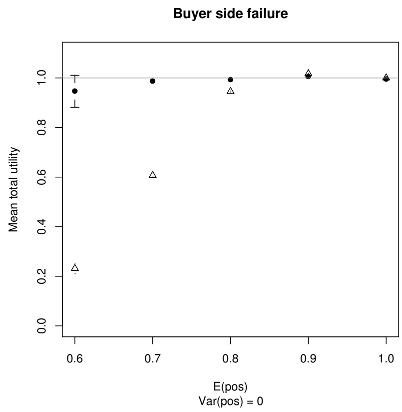

(c) Total utility

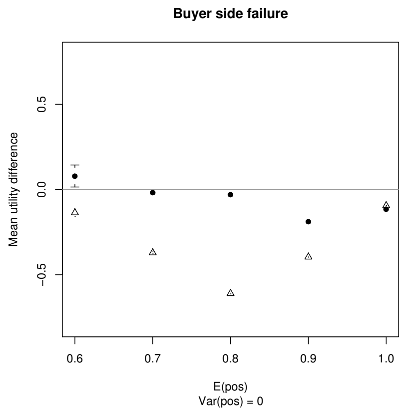

(d) Utility difference

Figure C.5: Buyer-side failure, with $Var(pos) = 0$, comparing T-ZIP (circles) and ZI (triangles).

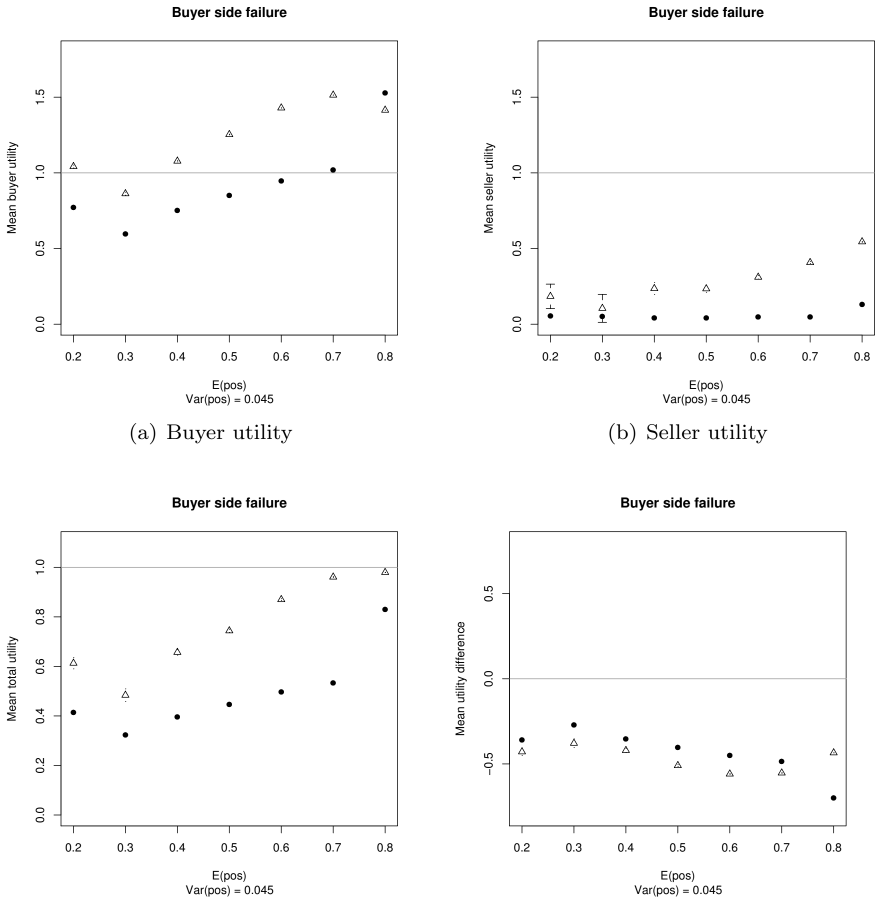

(a) Buyer utility

(b) Seller utility

(c) Total utility

(d) Utility difference

Figure C.6: Buyer-side failure, with $Var(pos) = 0.045$, comparing T-ZIP (circles) and ZI (triangles).

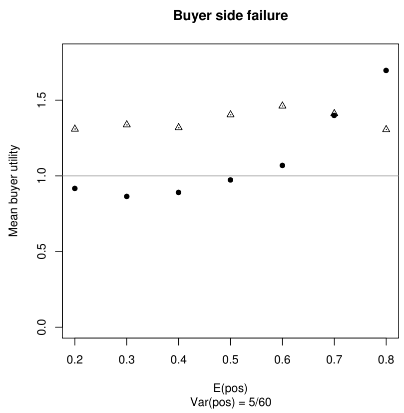

(a) Buyer utility

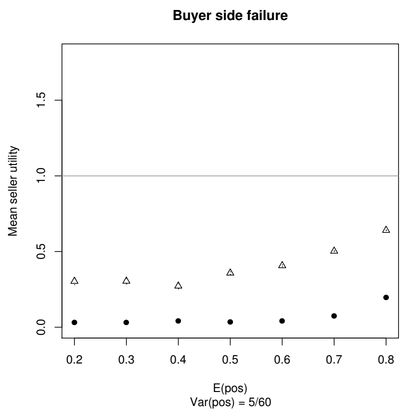

(b) Seller utility

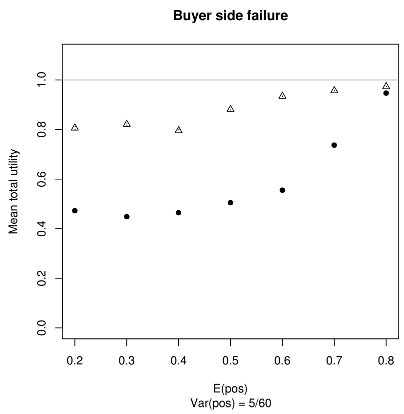

(c) Total utility

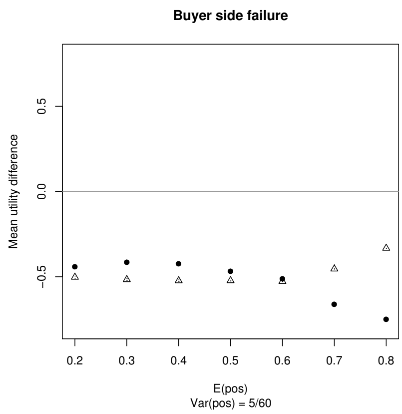

(d) Utility difference

Figure C.7: Buyer-side failure, with $Var(pos) = 0.083$, comparing T-ZIP (circles) and ZI (triangles).

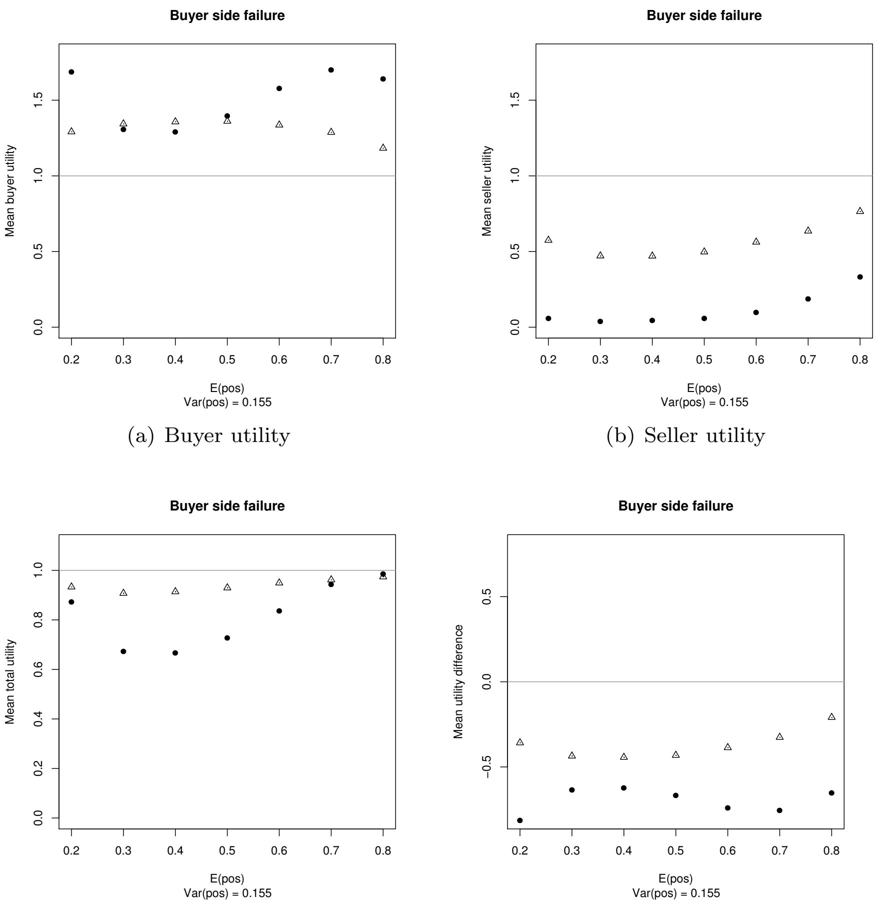

(a) Buyer utility

(b) Seller utility

(c) Total utility

(d) Utility difference

Figure C.8: Buyer-side failure, with $Var(pos) = 0.155$, comparing T-ZIP (circles) and ZI (triangles).

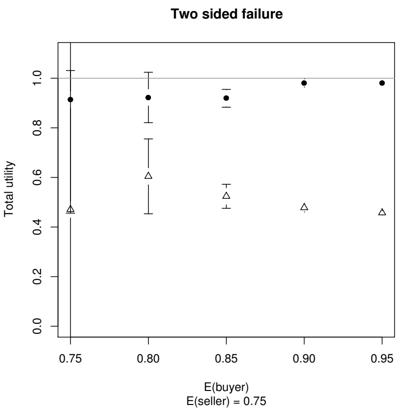

(a) $E(pos:S) = 0.75$

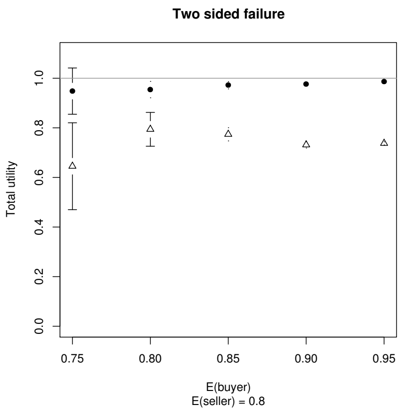

(b) $E(pos:S) = 0.8$

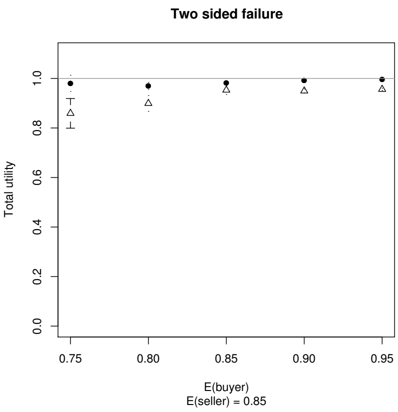

(c) $E(pos:S) = 0.85$

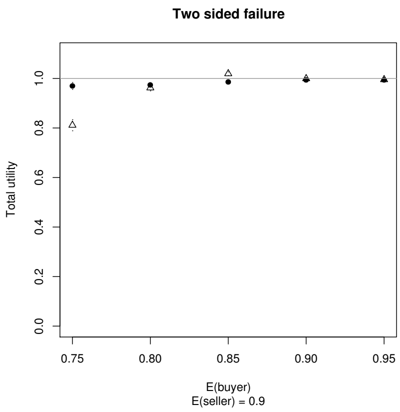

(d) $E(pos:S) = 0.9$

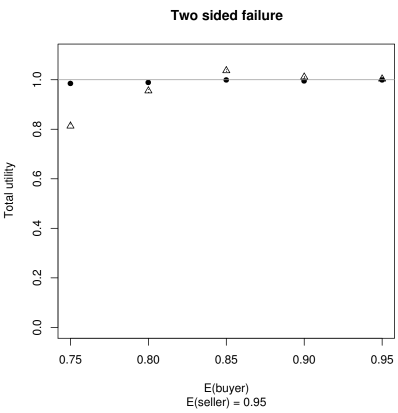

(e) $E(pos:S) = 0.95$

Figure C.9: Total utility with two-sided failure. Each subfigure shows a fixed level of POS for sellers, while the level of POS for buyers is varied.

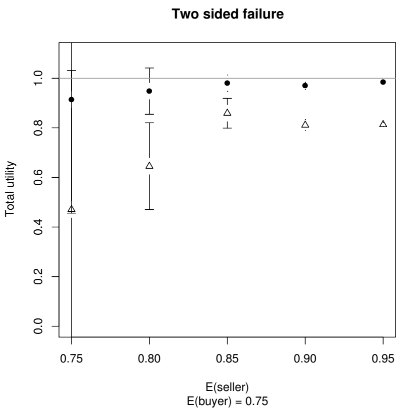

(a) $E(pos : B) = 0.75$

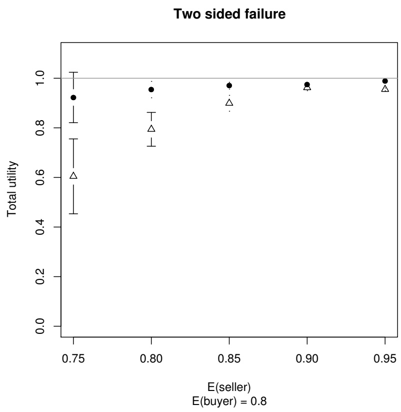

(b) $E(pos : B) = 0.8$

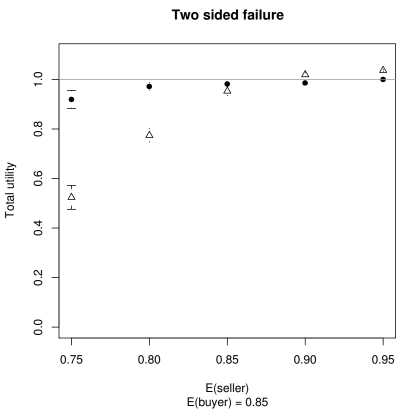

(c) $E(pos : B) = 0.85$

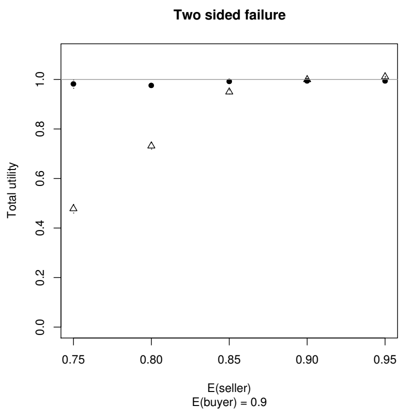

(d) $E(pos : B) = 0.9$

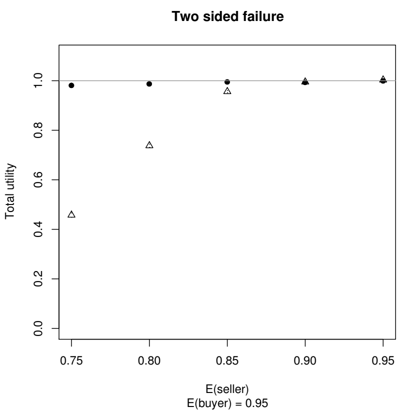

(e) $E(pos : B) = 0.95$

Figure C.10: Total utility with two-sided failure. Each subfigure shows a fixed level of POS for buyers, while the level of POS for sellers is varied.

Appendix D

T-ZIP: Instantaneous Response Step

In this appendix, figures are provided that enable comparison of the T-CDA with T-ZIP traders where the IRS is enabled and disabled. These results should be compared on a qualitative basis: although a great correspondence between the two conditions is expected, quantitatively identical results are not.

The first four figures enable comparison on the basis of price volatility, market efficiency and number of transactions over trading days. The final four figures show transaction time distributions for IRS enabled and IRS disabled runs side by side in a box plot. There are two sub-figures: one where times are plotted as-is, and one where the transaction times with IRS enabled are multiplied by twelve, to show the correspondence between the resulting boxes. The left-hand box has IRS disabled, whereas the right-hand box has IRS enabled.

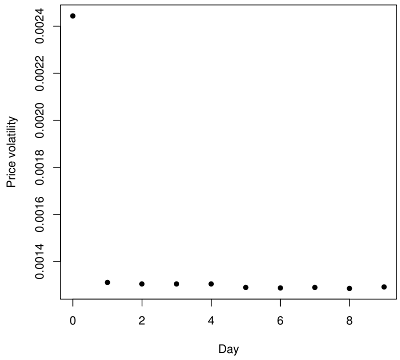

(a) Price volatility with IRS enabled

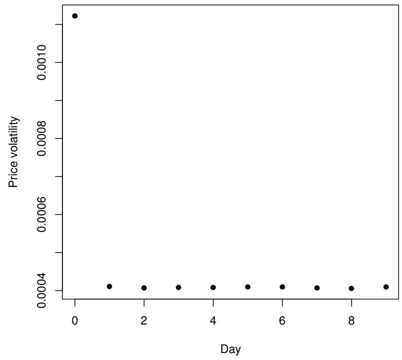

(b) Price volatility with IRS disabled

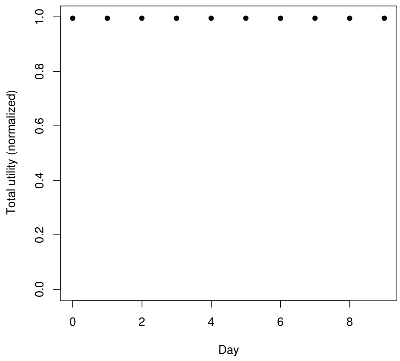

(c) Efficiency with IRS enabled

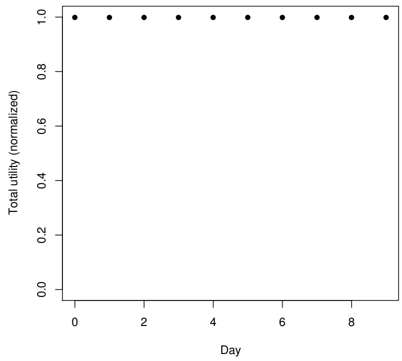

(d) Efficiency with IRS disabled

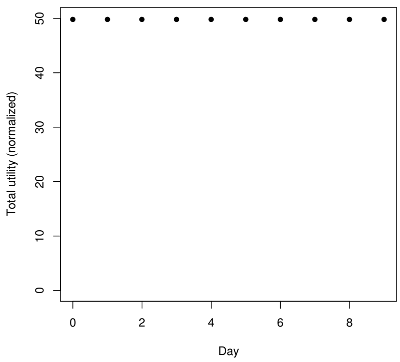

(e) Transactions with IRS enabled

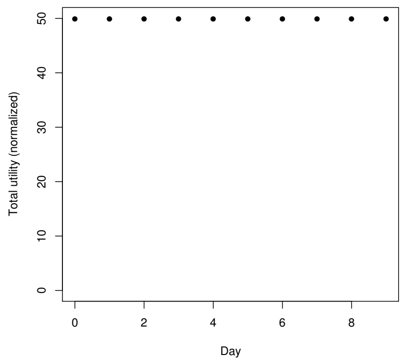

(f) Transactions with IRS disabled

Figure D.1: No execution uncertainty (CDA).

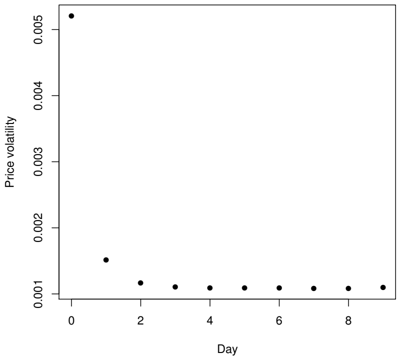

(a) Price volatility with IRS enabled

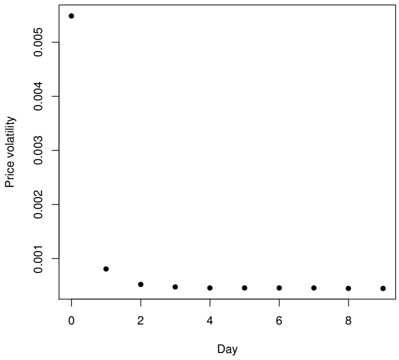

(b) Price volatility with IRS disabled

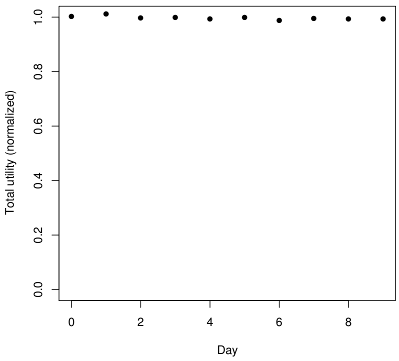

(c) Efficiency with IRS enabled

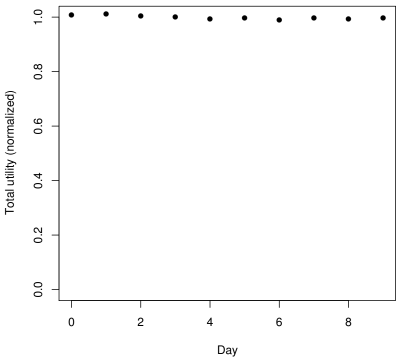

(d) Efficiency with IRS disabled

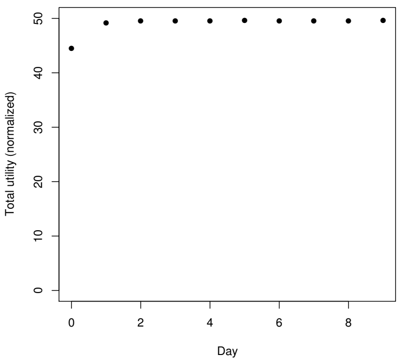

(e) Transactions with IRS enabled

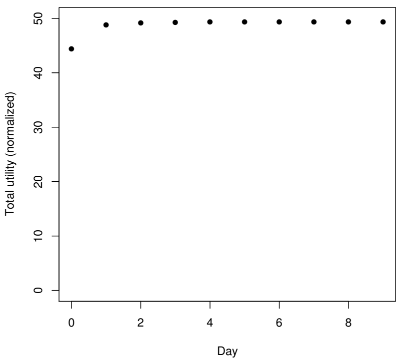

(f) Transactions with IRS disabled

Figure D.2: Buyer side failure, $E(pos) = 0.8$.

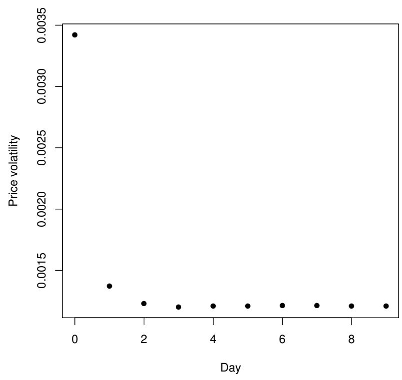

(a) Price volatility with IRS enabled

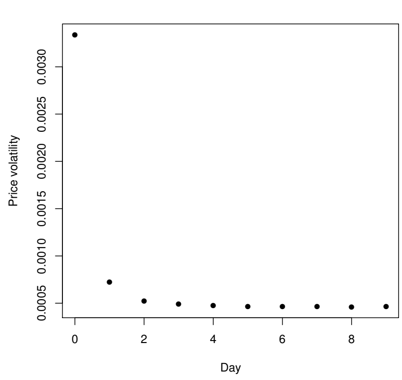

(b) Price volatility with IRS disabled

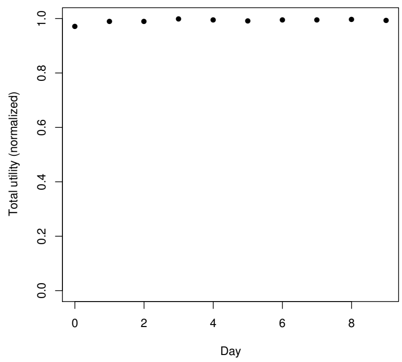

(c) Efficiency with IRS enabled

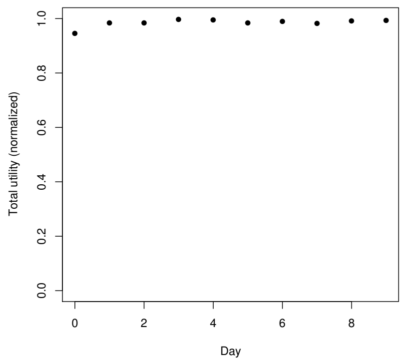

(d) Efficiency with IRS disabled

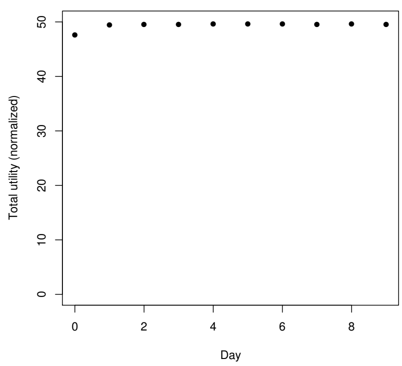

(e) Transactions with IRS enabled

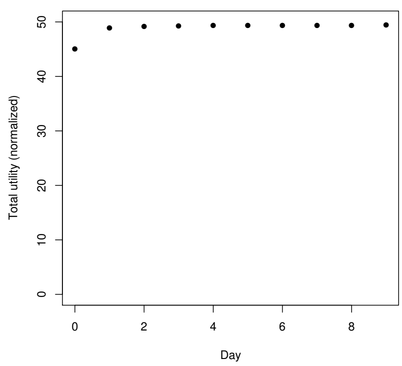

(f) Transactions with IRS disabled

Figure D.3: Seller side failure, $E(pos) = 0.8$.

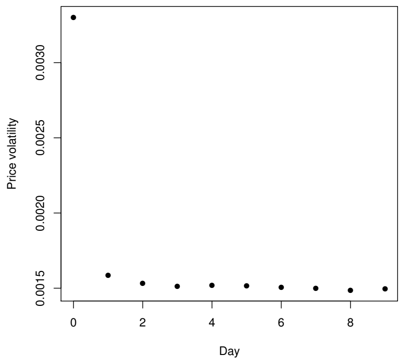

(a) Price volatility with IRS enabled

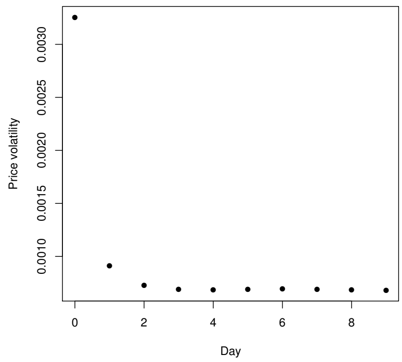

(b) Price volatility with IRS disabled

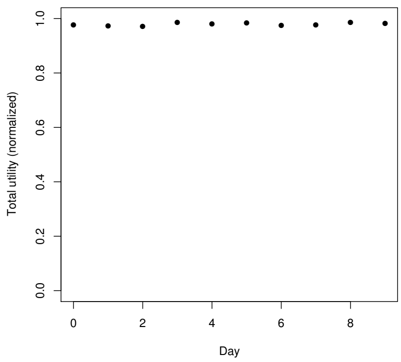

(c) Efficiency with IRS enabled

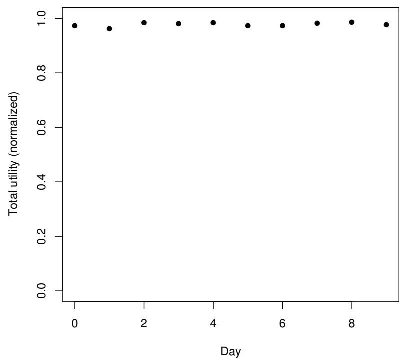

(d) Efficiency with IRS disabled

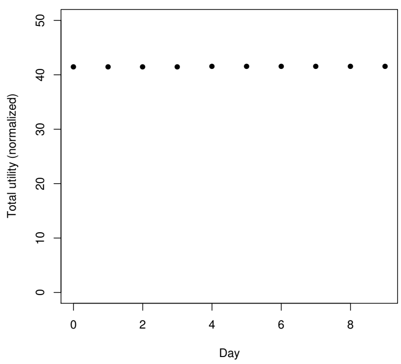

(e) Transactions with IRS enabled

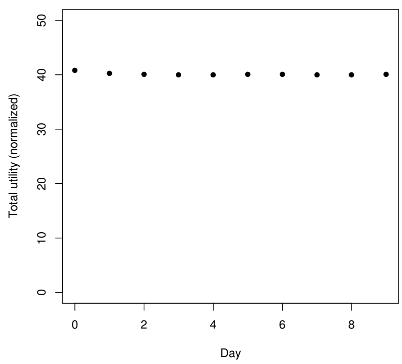

(f) Transactions with IRS disabled

Figure D.4: Two sided failure, $E(pos) = 0.85$.

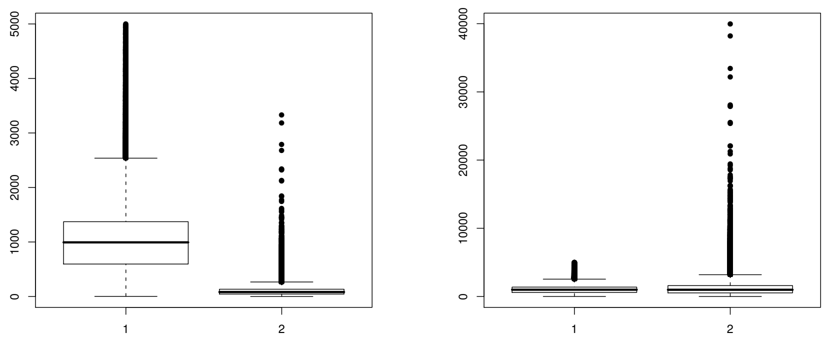

(a) Unscaled (b) Scaled (factor 12)

Figure D.5: Transaction timing in a CDA market.

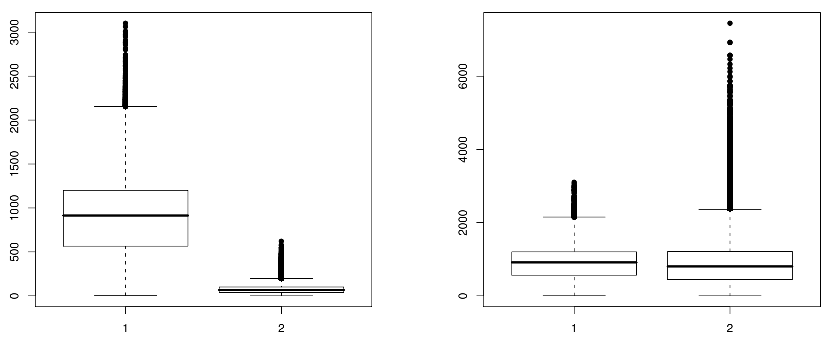

(a) Unscaled (b) Scaled (factor 12)

Figure D.6: Transaction timing in a buyer side failure market.

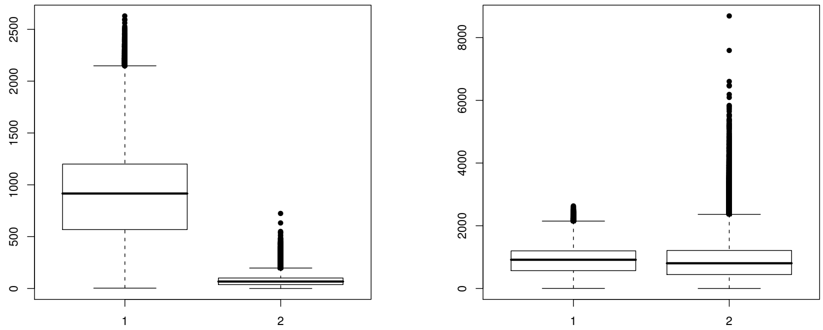

(a) Unscaled (b) Scaled (factor 12)

Figure D.7: Transaction timing in a seller side failure market.

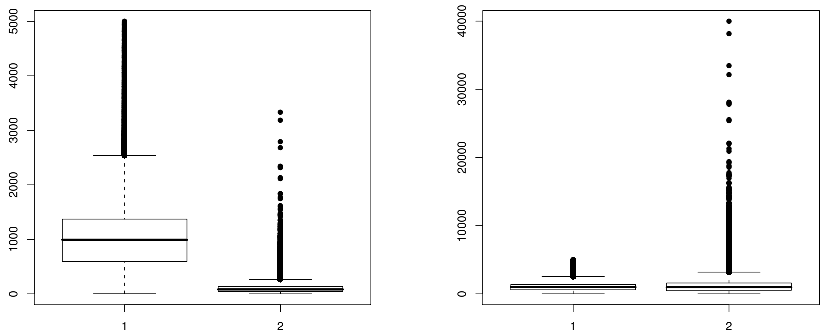

(a) Unscaled (b) Scaled (factor 12)

Figure D.8: Transaction timing in a two sided failure market.

Appendix E

The T-CDA Simulator

This appendix describes the simulation software that implements the T-CDA, the trading environment in which it operates and the agents and agent strategies that are evaluated in this thesis.

The T-CDA simulator is based on multi-unit CDA simulation software by Vytelingum (2006). The software has been extensively refactored to improve its overall design, robustness and configurability, as well as to accommodate the T-CDA mechanism. The graphical user interface (GUI), however, has been reused practically without modification. Figure E.1 provides a high-level overview of the architecture of the simulator.

Because the simulator was originally targeted at a different variant of the CDA and because one of the goals of refactoring was to preserve the original functionality, the simulator can, through configuration, support different market rules, agent behaviours and time models. The simulator can be run with GUI or without it. Data recording is pluggable (different modules can be selected based on experiment needs) and all parameter settings and selection of specific implementations (runner, auctioneer, strategies) is done through a plain-text settings file.

In the simulator, a market definition consists of the definition of several groups of traders (usually two, a group of buyers and a group of sellers). For each group, a strategy, endowment source, execution model and trust source is specified. The endowment source and execution model can be generated on-the-fly by the simulator, or read from a data file, to support repeatable experiments. The trust source is currently always generated on-the-fly.

The simulator can run any number of market definitions without restarting, or reloading its configuration. This is supported through the `cycle()` methods seen in Figure E.1, which move the simulation along to the next market definition. Moreover, the `reset()` methods reset the simulation to a clean state, so that one market definition can be run any number of times. Finally, the simulation can be (re-)started from any specific market definition, enabling a single simulation configuration to be run in parts, or restarted after a failure.

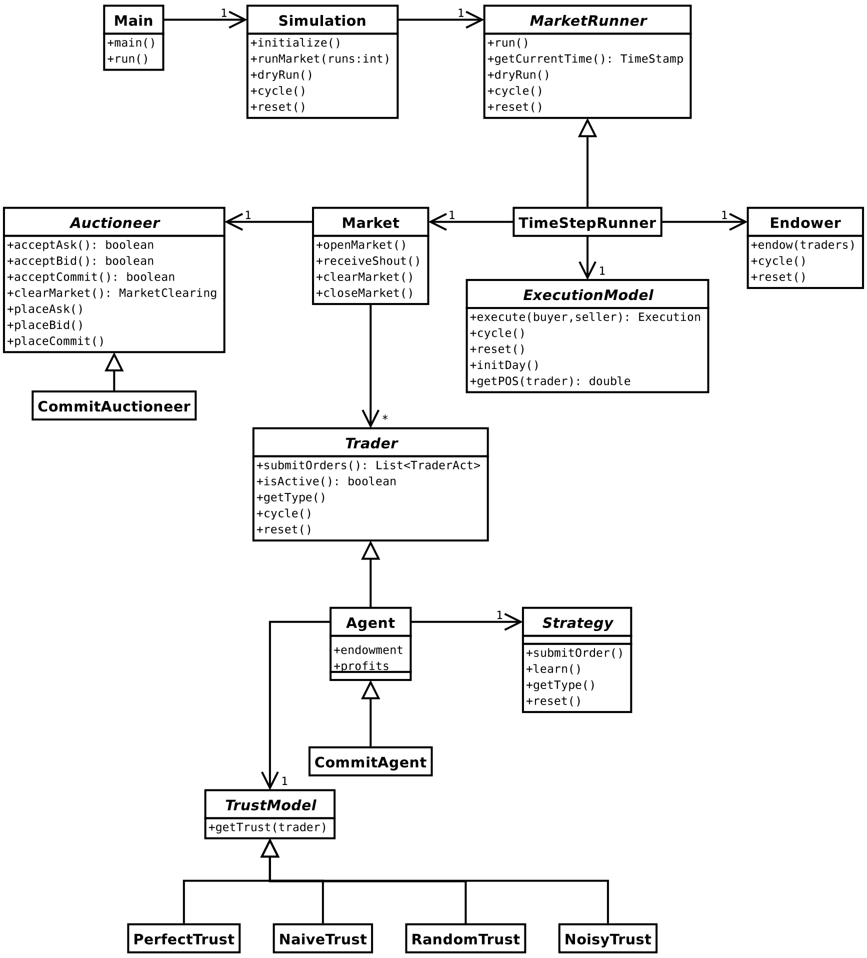

Figure E.1: Class diagram of the T-CDA simulator; provides a high-level overview only. The entry point is `Main.main()`. Note that through abstraction, different market rules (CDA, T-CDA and multi-unit CDA) and different agent behaviours as well as varied time models are supported. Only those parts relevant for the T-CDA are shown here, however.

References

AAMAS'02 (2002). *Proc 3rd Int Conf on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy.

AAMAS'04 (2004). *Proc 3rd Int Conf on Autonomous Agents and Multi-Agent Systems*, New York, USA.

Anderson, D. P., Cobb, J., Korpela, E., Lebofsky, M., and Werthimer, D. (2002). Seti@home: an experiment in public-resource computing. *Commun. ACM*, 45(11):56–61.

Buyya, R., Abramson, D., and Giddy, J. (2000). An economy driven resource management architecture for global computational power Grids. Las Vegas, USA. CSREA Press.

Buyya, R., Abramson, D., and Venugopal, S. (2005). The Grid economy. *Proc of the IEEE*, 93(3):698–714.

Buyya, R. and Bubendorfer, K. (2009). *Market-Oriented Grid and Utility Computing*. John Wiley & Sons Inc, New Jersey, USA.

Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-aguilar, J. A., and Sousa, P. (2006). Issues in multiagent resource allocation. *Informatica*, 30:2006.

Clearwater, S. H., editor (1995). *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, New Jersey, USA.

Cliff, D. (2005). Zip60: Further explorations in the evolutionary design of online auction market mechanisms. Technical Report HPL-2005-85, HP Laboratories Bristol.

Cliff, D. and Bruten, J. (1997). Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HPL-97-141, HP Laboratories Bristol.

Cohen, P. R. (1995). *Empirical Methods for Artificial Intelligence*. MIT Press.

Dash, R. K., Parkes, D. C., and Jennings, N. R. (2003). Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47.

Dash, R. K., Ramchurn, S. D., and Jennings, N. R. (2004). Trust-based mechanism design. In AAMAS'04 (2004), pages 748–755.

Dash, R. K., Vytelingum, P., Rogers, A., David, E., and Jennings, N. R. (2007). Market-based task allocation mechanisms for limited capacity suppliers. *IEEE Transactions on Systems, Man and Cybernetics (Part A)*, 37(3):391–405.

Foster, I., Jennings, N. R., and Kesselman, C. (2004). Brain meets brawn: Why Grid and agents need each other. In AAMAS'04 (2004), pages 8–15.

Foster, I. and Kesselman, C. (2003). *The Grid 2: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, San Francisco, CA, USA.

Foster, I., Kesselman, C., Tsudik, G., and Tuecke, S. (1998). A security architecture for computational grids. In *CCS '98: Proc of the 5th ACM Conference on Computer and Communications Security*, pages 83–92, New York, NY, USA. ACM.

Gjerstad, S. and Dickhaut, J. (1998). Price formation in double auctions. *Games and Economic Behavior*, 22:1–29.

Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *The Journal of Political Economy*, 101(1):119–137.

Gomoluch, J. and Schroeder, M. (2003). Market-based resource allocation for Grid computing: A model and simulation. In *Proc of the First International Workshop on Middleware for Grid Computing. Rio de*, pages 16–20.

He, M., Leung, H. F., and Jennings, N. R. (2003). A fuzzy logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Trans on Knowledge and Data Engineering*, 15:1345–1363.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Ogston, E. and Vassiliadis, S. (2002). A peer-to-peer agent auction. In AAMAS'02 (2002), pages 151–159.

Porter, R., Ronen, A., Shoham, Y., and Tennenholtz, M. (2008). Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799.

Pourebrahimi, B., Bertels, K., Kandru, G., and Vassiliadis, S. (2006). Market-based resource allocation in Grids. *e-Science and Grid Computing, 2006. e-Science '06. Second IEEE International Conference on*, pages 80–80.

Ramchurn, S. D., Huynh, T. D., and Jennings, N. R. (2004). Trust in multiagent systems. *The Knowledge Engineering Review*, 19(1):1–25.

Schnizler, B., Neumann, D., Veit, D., and Weinhardt, C. (2008). Trading grid services - a multi-attribute combinatorial approach. *European Journal of Operational Research*, 187(3):943 – 961.

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111.

Stein, S., Payne, T., and Jennings, N. (2008). Flexible service provisioning with advance agreements. In *Proc 7th Int Conf on Autonomous Agents and Multi-Agent Systems*, pages 249–256, Estoril, Portugal.

Tan, Z. and Gurd, J. R. (2007). Market-based grid resource allocation using a stable continuous double auction. In *Proc 8th IEEE/ACM Int Conf on Grid Computing*, pages 283–290, Austin TX, USA.

Teacy, W. T. L. (2006). *Agent-Based Trust and Reputation in the Context of Inaccurate Information Sources*. PhD thesis, University of Southampton.

Tesauro, G. and Bredin, J. L. (2002). Strategic sequential bidding in auctions using dynamic programming. In AAMAS'02 (2002), pages 591–598.

Vytelingum, P. (2006). *The Structure and Behaviour of the Continuous Double Auction*. PhD thesis, Southampton, UK.

Vytelingum, P., Cliff, D., and Jennings, N. (2008). Strategic bidding in continuous double auctions. *Artificial Intelligence Journal*, 172:1700–1729.

Wolski, R., Brevik, J., Plank, J. S., and Bryan, T. (2003). Grid resource allocation and control using computational economies. In Berman, F., Fox, G., and Hey, T., editors, *Grid Computing: Making The Global Infrastructure a Reality*, Wiley Series in Communications Networking & Distributed Systems, pages 747–771. John Wiley & Sons, Ltd.

Wolski, R., Plank, J. S., Brevik, J., and Bryan, T. (2001). Analyzing market-based resource allocation strategies for the computational Grid. *International Journal of High Performance Computing Applications*, 15:258–281.

Wooldridge, M. (2002). *An Introduction to Multi-Agent Systems*. J. Wiley.